Macroeconomic Shocks and Risk Premia

Gabor Pinter[*]

3rd August 2018

Comments Welcome

Abstract

What are the macroeconomic forces behind the cross-sectional and time-series variation in expected excess returns? To answer this question, this paper integrates models of empirical asset pricing with structural vector autoregressions (VAR). First, I use an unconditional asset pricing framework to construct an orthogonal shock in a macroeconomic VAR that best explains the cross-sectional variation in expected returns. The obtained "λ-shock" closely resembles identified monetary policy surprises and does not explain the recent US recessions. Second, I integrate return-forecasting methods to construct a second shock in the VAR, which best explains time-variation in expected returns. The obtained "γ-shock" turns out to be virtually orthogonal to the λ-shock, closely resembles demand-type financial shocks identified by macroeconomists, and explains most US recessions. I find that the λ-shock and the γ-shock jointly explain up to 80% of aggregate consumption fluctuations in the US.

JEL Classification*: C32, G12
Keywords: SDF, VAR, Shocks, Cross-section of Returns, Time-varying Risk Premia

1   Introduction

> *"In sum, we face two main questions. First, the equity premium question: What is there about recessions, or some other measure of economic bad times, that makes people particularly afraid that stocks will fall during those bad times—and so people require a large upfront premium to bear that risk? Second, the predictability question: What is there about recessions, or some other measure of economic bad times, that makes that premium rise—that makes people, in bad times, even more afraid of taking the same risk going forward? These are two separate questions..." (p. 947, Cochrane (2017))*

Understanding the macroeconomic forces behind the cross-sectional and time-series variation in expected excess asset returns continues to pose a challenge for both the macroeconomics and asset pricing literatures. Numerous variables have been proposed as pricing factors in cross-sectional asset pricing models or as predictors in return forecasting regressions, with the aim to better understand the drivers of risk premia. However, most of the variables proposed by the finance literature are reduced-form objects, which makes it difficult to establish robust relationships between structural macroeconomic forces and asset prices. In turn, the macroeconomics literature on structural vector autoregressions (VAR) has long been concerned with linking fluctuations in reduced-form variables to primitive aggregate shocks, leading to the development of numerous orthogonalisation techniques. However, most VAR models have ignored asset price information on the cross-sectional and time-series variation in expected returns, thereby remaining silent on how the identified macroeconomic shocks are related to the determination of risk premia. Reducing this gap between the finance and the macroeconomics literatures is therefore needed to get a better understanding of the empirical linkages between asset prices and economic fluctuations.

The objective of this paper is to address this challenge. To that end, and to answer the two questions in the opening quote, I propose an empirical framework that combines information both on the cross-sectional and time-series variation in expected returns with structural VAR techniques. I approach the problem from two different angles. First, I use the cross-section of asset prices, in a linear unconditional asset pricing framework with constant price of risk, to approximate innovations in the stochastic discount factor (SDF) with an orthogonal shock in a macroeconomic VAR model. This shock, which I refer to as the λ-shock, is constructed to explain the *cross-sectional variation* in expected returns. Second, I use the time-series variation in expected returns and look for an orthogonal shock in the VAR which drives the fluctuations in the macroeconomic variables that have the highest predictive power of future excess returns. This shock, which I refer to as the γ-shock, is constructed to explain *time-variation* in expected returns. I then study the dynamic effects of these shocks on the macroeconomy by making full use of the standard VAR toolkit, e.g. by inspecting impulse response functions (IRFs), forecast error variance (FEV) decomposition, historical decomposition and the estimated time-series of the shocks.

To sum up the main feature of the framework: instead of starting with a model of business cycles and testing its asset price implications, I work "backwards" by first using asset prices to construct the SDF from macroeconomic innovations and only *then* study its empirical relation to business cycles and to structural shocks identified by macroeconomists. By taking this reverse direction, one can circumvent issues that have long hindered the study of the empirical linkages between the macroeconomy and asset prices (consumption measurement errors, specification of utility functions, identification of structural shocks etc.).

Using a standard macroeconomic VAR and benchmark test portfolios for the US, my paper yields four main empirical results. First, I find that the λ-shock generates a negative comovement between the short-term interest rate and consumption, with the latter exhibiting a delayed response, consistent with the recent consumption-CAPM literature. Importantly, the λ-shock closely resembles monetary policy surprises as identified by the macroeconomics literature, often using very different methodologies (Romer and Romer 2004; Sims and Zha 2006; Gertler and Karadi 2015). This result highlights the overlap between linear pricing models of the cross-section of average returns (Fama and French, 1993) and structural shocks identified by the macroeconometric literature (Sims, 1980).

Second, the estimated γ-shock induces a sharp response in consumption and a positive comovement between the short-term interest rate and aggregate quantities. I find that the economic characteristics make the γ-shock resemble demand-type shocks as identified by the recent macroeconomic literature. Specifically, Christiano, Motto, and Rostagno (2014) used a linearised structural equilibrium model with financial frictions to show that exogenously fluctuating uncertainty related to the cross-section of risk ("risk shocks") can explain around 60% of US business cycle fluctuations. Though I use an empirical model and rely solely on information regarding time-variation in risk premia, I find a close empirical relationship between risk shocks and the estimated γ-shock. This result highlights the overlap between recent explanations of business cycle fluctuations, offered by the macroeconomic literature, and the drivers of time-varying risk premia, long studied by the asset pricing literature.

Third, even though the λ-shock and the γ-shock are not restricted to be orthogonal to each other, but a key empirical finding of this paper is that they are close to being orthogonal in the data. This implies that the stochastic macroeconomic drivers of average risk premia are empirically orthogonal to the stochastic drivers of variation in risk premia. I will discuss the link of this result with consumption-based asset pricing theories.

Fourth, given the orthogonality of the λ-shock to the γ-shock in the data, I compute FEV and historical decompositions to assess the contribution of these shocks to the US business cycle. I find that the γ-shock explains most US recessions in my sample and most of the high-frequency variation in aggregate consumption. In turn, the λ-shock is important in driving lower-frequency variation in aggregate consumption, and it made a large contribution only to the recession in the early 1980s. Importantly, while the construction of the λ-shock and the γ-shock relies exclusively on asset price information and not on macroeconomic assumptions,

I find that these two shocks jointly explain up to 80% of aggregate consumption fluctuations over the 1963-2015 period.

The (cross-sectional) method to construct the λ-shock connects two simple ideas: (i) a basic fact of the empirical finance literature (Cochrane, 2005) is that β-pricing models of the cross-section of asset prices imply a linear model of the stochastic discount factor (SDF); (ii) a basic fact of the macroeconometrics literature (Sims, 1980) is that orthogonalised shocks in a VAR model are linear combinations of the reduced-form innovations. These two facts imply that, given the space spanned by the innovations of a linear VAR and the space spanned by the cross-section of asset returns, one can construct orthogonal shocks in the VAR that are best linear approximations of the SDF (with all other shocks in the VAR demanding zero average risk premia).

The (time-series) method to construct the γ-shock builds on the asset pricing literature which found empirical evidence on the predictability of excess returns by financial and macroeconomic variables, implying that expected excess returns vary with the business cycle (Cochrane 2011). This literature typically employed univariate time-series techniques to regress realised excess returns on lagged values of valuation ratios (Campbell and Shiller 1988; Fama and French 1988) or macroeconomic variables (Fama and French 1989; Ferson and Harvey 1991; Lettau and Ludvigson 2001a), and assessed the forecasting power of the proposed predictors based on the regression R^2 statistic. Given that most of the proposed predictor variables are reduced-form objects, their forecasting power could in theory be decomposed to the historical contribution of primitive economic shocks that generated fluctuations in the given predictors. I take this idea to the limit, and search for a single orthogonal shock in my macroeconomic VAR with the following property: the historical contribution of this shock to predictor variables in the VAR would generate counterfactual variation in these predictors, which would have the highest possible R^2 statistic when using them in return forecasting regressions. To the extent that time-variation in expected returns is linked to economic booms and busts (Lettau and Ludvigson, 2010; Cochrane, 2011), the γ-shock can be thought of as the stochastic driver of recessions in the VAR.

It is to note that constructing the λ-shock and the γ-shock does not classify as identification in the classical macroeconometric sense. The methods do not build on assumptions typically used in VARs to identify structural primitives such as aggregate supply or demand shocks. One interpretation of these shocks is via the generalisation of the discount factor (M_{t+1}) implied by consumption-based asset pricing:

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} Y_{t+1}, \tag{1.1}$$

where C_t is consumption of the representative household, β is the subjective discount rate, σ is the risk aversion coefficient, and Y_{t+1} is a key state variable, directly related to recessions and to time-varying risk-bearing ability, as discussed in Cochrane (2017). Inspecting through

the lens of this framework, as will be discussed in further detail, the λ-shock can be thought of as innovations in the consumption growth process ($\frac{C_{t+1}}{C_t}$) which explain the average level of expected returns; the γ-shock can be thought of as innovations in the recession-related state variable (Y_{t+1}) which explain time-variation in expected returns. Numerous papers in both the asset pricing and macroeconomics literatures have provided theoretical explanations of the drivers of the discount factor (1.1).[1] A key advantage of my framework is its reverse engineering nature: it provides a simple and agnostic way to estimate the stochastic drivers of the discount factor, and their effects on the business cycle, without having to specify the functional form of the household's utility function (1.1) and the corresponding structural model. I will show that this reverse direction may be preferred to the more obvious direction – using identified structural shocks directly in asset pricing tests.

In addition, modelling the macroeconomic dynamics as a VAR (the space from which I recover the λ-shock) separately from the cross-section of asset prices (the space that induces the λ-shock) provides further applications. One can explore how different test portfolios may proxy different macroeconomic risks in a linear unconditional asset pricing framework. In the proposed framework, this can be done by changing the test assets while keeping the state-variables in the VAR fixed and merely rotating the reduced-form VAR innovations.[2] I find that the impulse response functions of the λ-shock is similar across standard equity portfolios.[3] Importantly, the λ-shock implied by government bond returns is similar to that implied by equities as well. This evidence adds to the growing literature on the joint pricing of stocks and bonds (Lettau and Wachter 2011; Bryzgalova and Julliard 2015; Adrian, Crump, and Moench 2015; Koijen, Lustig, and Van Nieuwerburgh 2017).

An inverse exercise is to fix the test assets (thereby controlling for the proxied aggregate risk) and change the state variables in the VAR. This can serve two purposes. First, given the proliferation of asset pricing factors in the finance literature (Harvey, Liu, and Zhu, 2016), the proposed VAR framework could model the joint dynamics of any reduced-form variables that individually have been found to price the cross-section of returns, and to link the common stochastic driver of these variables to a single or multiple orthogonal shocks.

---

[1] In general equilibrium models, used to study risk premia, the λ-shock was traditionally linked to technology shocks (Mehra and Prescott 1985; Jermann 1998). The γ-shock in the asset pricing literature could represent shocks to the volatility of the consumption process (Bansal and Yaron, 2004) or shocks to time-varying cross-sectional variance of individual consumption growth (Constantinides and Duffie, 1996), amongst other alternative theories. In the macroeconomics literature, the γ-shock typically corresponds to "preference shocks" that were found to be important contributors to short-term business cycle fluctuations in quantitative New Keynesian models (Smets and Wouters, 2007).

[2] This way I control for the macroeconomic information set when changing the test assets which presents an advantage over standard no-arbitrage estimation of the SDF using observed prices. As explained in the next Section, my method exploits the fact that structural VARs are not identified, i.e. there are infinite combinations of possible orthogonalisations that conform to the reduced-form variance-covariance matrix of the VAR. This is the degree of freedom which allows me to estimate a different λ-shock each time I change the test portfolios without having to change the space that I recover the estimates of λ-shock from. See Section A.3 of the Appendix for further discussion and Figure 7 for a pictorial illustration.

[3] This is in spite of the fact that changing test assets can substantially change the pricing performance of the given VAR model, as discussed in Section B.9 of the Appendix.

Second, one can explore how the realisation of aggregate risks proxied by the given test assets may affect different parts of the macroeconomy. To illustrate both these points, I will add to the VAR the aggregate capital ratio of the financial intermediary sector, constructed by He, Kelly, and Manela (2017): this reduced-form variable[4] has been successful in explaining the cross-section of excess returns; and the Great Recession sparked increasing interest in the role of financial intermediaries in affecting risk premia (Adrian, Etula, and Muir, 2014) and the business cycle in general (Gertler and Kiyotaki 2010; Brunnermeier and Sannikov 2014; He and Krishnamurthy 2014). I find that about two-third of the variation in financial intermediary capital can be explained jointly by the λ-shock and the γ-shock. Importantly, the contributions of both shocks (inducing very different macroeconomic dynamics) to the historical volatility of the capital ratio are about equally important. This highlights that by purely focusing on the reduced-form unforecastable component in the capital ratio, one cannot accurately detect the nature of the macroeconomic force responsible for the observed fluctuations in intermediary balance sheets and their implications for asset prices and the wider economy.

A number of additional robustness exercises will be included in the paper. Importantly, I propose a third orthogonalisation scheme (in addition to constructing the λ-shock and the γ-shock), which is explicitly designed to explain recessions without any reference to time-variation in expected returns. This shock, which I refer to as a Recession-shock, uses an agnostic identification scheme based on the historical decomposition of aggregate consumption growth during the Great Recession. The identifying assumption is that the historical contribution of the Recession-shock explains the collapse of consumption over the 2008Q1-2009Q4 period. Given that the countercyclical relationship between recessions and risk premia is by now an established sylised macro-finance fact (Cochrane, 2011, 2017), this third method should deliver a shock which behaves very similarly to the γ-shock, whose construction is solely based on variation in excess market return, without any reference to recessions. My empirical evidence indeed confirms this.

Further, I check whether my results are robust to countries other than the US. I apply the proposed methodology to UK data, and find that the behaviour of the λ-shock and the γ-shock is similar across the two countries.

Related Literature   My paper is related to the literature on finding macroeconomic factors that drive the cross-sectional and time-series variation in expected returns. A partial list includes Chen, Roll, and Ross (1986), Fama and French (1989), Ferson and Harvey (1991), Campbell (1996), Cochrane (1996), Lettau and Ludvigson (2001a,b), Petkova (2006), Piazzesi, Schneider, and Tuzel (2007), Liu and Zhang (2008), Maio and Santa-Clara

---

[4]A simple illustration of this fact is that the intermediary capital risk factor and excess returns on the market (the second pricing factor in He, Kelly, and Manela (2017)'s two-factor model) have a 78% correlation with one another.

(2012), Boons and Tamoni (2015), He, Kelly, and Manela (2017) and Koijen, Lustig, and Van Nieuwerburgh (2017) amongst others. My paper also relates to the empirical literature on the link between monetary policy and expected returns (Gurkaynak, Sack, and Swanson 2005; Bernanke and Kuttner 2005; Lucca and Moench 2015; Weber 2015; Chava and Hsu 2015; Ozdagli and Velikov 2016).

The method to construct the λ-shock draws on the structural VAR literature that uses sign restrictions to identify structural shocks (Uhlig 2005; Rubio-Ramirez, Waggoner, and Zha 2010; Fry and Pagan 2011; Baumeister and Hamilton 2015). The method to construct the γ-shock draws on the macroeconomic literature on the identification of news shocks (Uhlig 2004; Barsky and Sims 2011; Kurmann and Otrok 2013); I combine the ideas from this literature with predictive regressions of the asset pricing literature (Campbell and Shiller 1988; Goyal and Welch 2008; Pastor and Stambaugh 2009). The method to identify a Recession-shock based on matching historical decompositions is inspired by papers matching impulse response functions (Christiano, Eichenbaum, and Evans, 2005) and the recent structural VAR literature (Ludvigson, Ma, and Ng, 2017; Antolin-Diaz and Rubio-Ramirez, 2018; Ben Zeev, 2018).

The remainder of the paper is as follows: Section 2 explains the method, Section 3 presents the empirical results, Section 4 provides robustness checks and extensions, and Section 5 concludes.

2 The Econometric Framework

2.1 The λ-Shock

Assume that the macroeconomy evolves according to a k-variable reduced-form VAR:

$$X_t = c + \mathbb{A}_1 X_{t-1} + \cdots + \mathbb{A}_p X_{t-p} + \eta_t, \qquad \eta_t \sim N(0, \Omega), \tag{2.1}$$

where the reduced form innovations, η_t, are related to the structural shocks e_t by an invertible matrix $\boldsymbol{B}$, $\eta_t = \boldsymbol{B} e_t$. Following Uhlig (2005), I refer to the columns of $\boldsymbol{B}$ as *impulse vectors*.

Finance papers using VARs (Campbell, 1996; Petkova, 2006; Hansen, Heaton, and Li, 2008; Boons, 2016) often used Cholesky decomposition ($\boldsymbol{B} = chol(\Omega)$) to obtain orthonormalised shocks as pricing factors as estimates of sources of aggregate risks.[5] Building on these papers, I explore the whole space of possible orthonormalisations to approximate the SDF from linear combinations of residuals η_t. Specifically, I select a single or

---

[5]Originally, Sims (1980) applied Cholesky decomposition to obtain a triangular structure in the spirit of Wold (1954). A plethora of new techniques have been proposed by the macroeconometrics literature to provide full or partial identification of $\boldsymbol{B}$, involving both point and set identification of the elements of $\boldsymbol{B}$. See Kilian and Lutkepohl (2016), Ramey (2016), and Ludvigson, Ma, and Ng (2017) for a recent review of the literature.

multiple impulse vectors such that the shocks associated with all other impulse vectors are *orthogonal* to the SDF. To implement the method, I first estimate the reduced-form VAR (2.1) and apply Cholesky decomposition to the estimated variance-covariance matrix $\hat{\boldsymbol{B}} = chol\left(\hat{\Omega}\right)$. One can take any orthonormal matrix Q to obtain a new structural impact matrix $\hat{\boldsymbol{B}}^{\star} = \hat{\boldsymbol{B}}Q$, thereby obtaining a new set of orthogonal shocks, which conforms to $\hat{\Omega}$, i.e. $\hat{\Omega} = \hat{\boldsymbol{B}}^{\star}\left(\hat{\boldsymbol{B}}^{\star}\right)' = \hat{\boldsymbol{B}}Q\left(\hat{\boldsymbol{B}}Q\right)' = \hat{\boldsymbol{B}}\hat{\boldsymbol{B}}'$. One could think of Q as a rotation matrix, e.g. $Q = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ with the associated Euler-angle θ. The next proposition for the two-dimensional $\mathbb{R}^2$ case highlights how to find the rotation which generates the λ-shock.

Proposition 1 *Given the linear combination: $m = \lambda_1 f_1 + \lambda_2 f_2$, where $\lambda_1, \lambda_2 \in \mathbb{R}$, $m, f_1, f_2 \in \mathbb{R}^2$, $\|f_1\| = \|f_2\| = 1$ and $\langle f_1, f_2 \rangle = 0$, $\exists$ matrix $Q = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ such that $m = \lambda_1^{\star} f_1^{\star} + \lambda_2^{\star} f_2^{\star}$, where $\lambda_1^{\star} = \|m\| \neq 0$, $\lambda_2^{\star} = 0$ and $f_i^{\star} = Q f_i$ for $i = 1, 2$.*

All proofs are in Appendix A. Figure 1 provides a graphical illustration via an example, whereby a linear model $m = 2f_1 + f_2$ with $f_1 \perp f_2$ is transformed to $m = \|m\| f_1^{\star} + 0 \cdot f_2^{\star}$ with $Q = \begin{bmatrix} \cos\theta^{\star} & -\sin\theta^{\star} \\ \sin\theta^{\star} & \cos\theta^{\star} \end{bmatrix}$, $\theta^{\star} = arctan\left(\frac{1}{2}\right)$, $\|f_1\| = \|f_2\| = \|f_1^{\star}\| = \|f_2^{\star}\| = 1$ and $\|m\| = \sqrt{5}$.

Figure 1: Graphical Illustration of Constructing the λ-shock
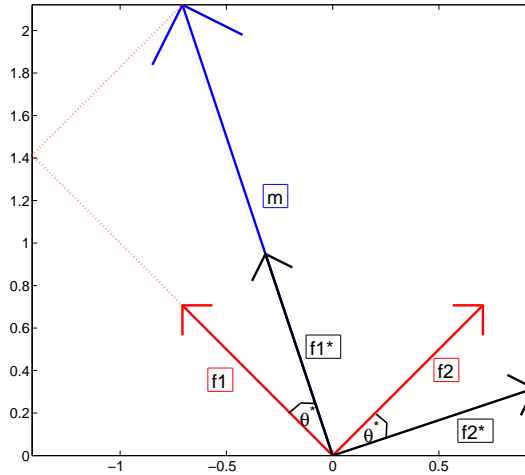
Notes: the 2-dimensional coordinate system illustrates the space spanned by reduced-form VAR innovations. The space contains vector m, which is the best linear approximation of the SDF, according to some test assets. The red perpendicular arrows ($f_1 \perp f_2$) illustrate an arbitrary orthonormalisation of the reduced-form VAR innovations, e.g. Cholesky decomposition as in Campbell (1996), Petkova (2006), Boons (2016) amongst others. Rotation and orthonormalisation do not change the spanned space, thereby leaving the information set and m unchanged. Therefore there exists $\theta^{\star}$ such that $m = \|m\| f_1^{\star}$ and $m \perp f_2^{\star}$.

While the proposition is a trivial piece of linear algebra, it has important implications for using orthonormalised shocks from VARs as pricing factors in linear pricing models. Given

that β-pricing models are equivalent to linear models of the SDF (Section A.2 of the Appendix), finding the Euler-angle θ and the associated rotation in a VAR (of *any* dimension) that delivers an orthonormalised shock with the highest price of risk ($\sqrt{5}$ and $f_1^\star$ in Figure 1) when pricing given test assets is equivalent to finding the best linear approximation of the SDF that lies in the innovation space of the VAR. By construction, all other orthonormalised shocks in the VAR ($f_2^\star$ in Figure 1) will be orthogonal the implied SDF and demand zero risk premia. Importantly, one can apply structural VAR tools to the obtained Euler-angles to study the link between the shock and macroeconomic dynamics. Given the VAR model, the rotation Q naturally depends on the test assets that induce the SDF. Section A.3 of the Appendix provides further illustration and highlights the geometric nature of the idea.

To find a $k \times k$ Q matrix in the VAR model 2.1, I span the space with Givens rotations to construct an orthonormal shock such that, given the test assets, the corresponding $\hat{\lambda}$ estimate in the second-pass Fama and MacBeth (1973) regression is maximised.[6] Example (2) illustrates the mechanics.

Example 2 (A Bivariate VAR Model) *Let R be a $T \times n$ matrix of excess returns of n test portfolios. Take a two-variable VAR model ($k = 2$) where the pricing factors are orthonormalised shocks given by Cholesky decomposition ($f_t = [f_{1t} \vdots f_{2t}] = \eta_t \boldsymbol{B}^{-1} = \eta_t \left(chol\left(\Omega\right)\right)^{-1}$). The implied model of the SDF (m) is:*

$$m_t = \lambda_1 f_{1t} + \lambda_2 f_{2t}, \tag{2.2}$$

where λ_1 and λ_2 are the prices of risk associated with f_1 and f_2. Given that the factors are not persistent (Adrian, Crump, and Moench, 2015), the λs can be estimated with the two-stage procedure of Fama and MacBeth (1973).[7] Because $f_1 \perp f_2$ and $var(f_1) = var\left(f_2\right) = 1$, the variance of the SDF is simply the sum of the squared values of the estimated prices of risk associated with the two shocks:

$$var\left(\hat{m}_t\right) = \hat{\lambda}_1^2 + \hat{\lambda}_2^2. \tag{2.3}$$

Rotation does not change the information set: the volatility of the implied SDF is determined by the specification of the VAR and not by rotating the variance-covariance matrix of the residuals. The main implication of proposition 1 is that the information in the VAR residuals can be summarised by only one orthogonal shock after an appropriate rotation, i.e. there exists matrix $Q = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ such that using $f_i^\star = Q f_i$ for $i = 1, 2$ as pricing factors,

---

[6]See Section A.5 of the Appendix and the sign restrictions literature (Uhlig 2005; Rubio-Ramirez, Waggoner, and Zha 2010; Fry and Pagan 2011; Kilian and Murphy 2012; Baumeister and Hamilton 2015) for more information on Givens rotations and QR decompositions.

[7]First, estimate n time series regressions, $R_{it} = a_i + f_t \beta_i + \epsilon_{it}$, $i = 1 \ldots n$. Then, estimate a cross-section regression, $\bar{R}_i = \tilde{\beta}_i \times \lambda + \alpha_i$, where $\bar{R}_i = \frac{1}{T} \sum_{t=1}^{T} R_{it}$, $\tilde{\beta}_i$ is the OLS estimate obtained in the first stage and α_i is a pricing error.

one of the estimated prices of risk is $\lambda_1^\star = \sqrt{var(\hat{m})}$, as the other one is zero, $\lambda_2^\star = 0$. This implies that the best approximation of the SDF is found, $f_1^\star = \hat{m}$, and Q can be used to perform structural analysis in the VAR, i.e. $\hat{\boldsymbol{B}}^\star = \hat{\boldsymbol{B}}Q$ can be used to compute IRFs.[8]

A linear model of the SDF, that uses arbitrarily orthonormalised VAR residuals, uniquely pins down one of the rows of the matrix $\left(\hat{\boldsymbol{B}}^\star\right)^{-1}$. However, this is not sufficient to carry out structural VAR analysis, because to do so one needs to know the column in the structural impact matrix $\hat{\boldsymbol{B}}^\star$ (and not its inverse $\left(\hat{\boldsymbol{B}}^\star\right)^{-1}$). Building on Example (2), the following proposition establishes the relationship between the angle θ needed to compute the impulse vector associated with the λ-shock.

Proposition 3 *Given the VAR model in Example (2) with variance-covariance matrix* $\Omega = \begin{bmatrix} \omega_{11} & \omega_{12} \\ \omega_{12} & \omega_{22} \end{bmatrix}$, *the column of* $\boldsymbol{B}^\star = \begin{bmatrix} B_{11}^\star & B_{12}^\star \\ B_{21}^\star & B_{22}^\star \end{bmatrix}$ *corresponding to the contemporaneous effect of the λ-shock is given by:* $B_{11}^\star = \sqrt{\omega_{11}}\cos(\theta)$ *and* $B_{21}^\star = \frac{\omega_{12}}{\sqrt{\omega_{11}}}\cos(\theta) + \sqrt{\omega_{22} - \omega_{12}^2/\omega_{11}}\sin(\theta)$, *with the rotation angle θ given by the prices of risks,* $\theta = \arcsin\left(\lambda_2/\sqrt{(\lambda_1^2 + \lambda_2^2)}\right)$.

As an empirical illustration of this statement, consider the following example.

Example 4 (A Bivariate VAR and the Consumption-CAPM) *Let the two variables in a quarterly VAR(2) be the log of consumption and the term spread, and let the test assets be the 25 Fama-French (FF25) portfolios. An OLS regression using data from 1970Q1 to 2012Q2 yields an estimated variance-covariance matrix* $\hat{\Omega} = \begin{bmatrix} 0.38 & -0.07 \\ -0.07 & 0.78 \end{bmatrix}$. *Given the reduced-form residuals, u_t^C and u_t^{Term}, the implied linear model of the SDF is $\hat{m}_t = 0.21 u_t^C + 1.14 u_t^{Term}$, implying that term-spread innovations load more on the SDF than consumption innovations. Consider the contemporaneous impact of the shock on the variables. Using the appropriate angle θ, the elements of the $\hat{B}^\star$ are:*

$$\begin{aligned} \hat{B}_{11}^\star = 0.006 \quad & \hat{B}_{12}^\star = 0.61 \\ \hat{B}_{21}^\star = 0.88 \quad & \hat{B}_{22}^\star = -0.12. \end{aligned} \tag{2.4}$$

The first column shows that a s.d. λ-shock induces a large (0.88pp) jump in the term spread, but has virtually no contemporaneous effect ($<0.01\%$) on consumption. The second column shows the contemporaneous effect of the shock that is by construction orthogonal to the implied SDF, thus demanding zero risk premia. This shock has a large (0.61%) contemporaneous effect on consumption which implies that virtually all of the one period ahead FEV in consumption is explained by a shock, exposure to which demands zero risk compensation according to the FF25.

---

[8]It is worth noting that finding Q is not needed to find the time-series λ-shock. Applying the Fama and MacBeth (1973) procedure to any linear combinations of the VAR innovations will produce a unique time-series of the λ-shock that can be obtained as the fitted values of the second-stage regression. This is highlighted by lemma 8 of the Appendix.

This simple example highlights the empirical relevance of the Consumption-CAPM literature which emphasises that news about current consumption growth are irrelevant to determining the level of risk premia.[9] To estimate the effect of the λ-shock on consumption dynamics at longer horizons, one needs to compute impulse response functions.

Definition 5 (Impulse Response Functions) *Consider a VAR(2) model $X_t = c + \mathbb{A}_1 X_{t-1} + \mathbb{A}_2 X_{t-2} + \eta_t$ with response matrices $\Phi_0 = I$, $\Phi_1 = \Phi_0 \mathbb{A}_1$, $\Phi_2 = \Phi_1 \mathbb{A}_1 + \Phi_0 \mathbb{A}_2$, $\ldots$, $\Phi_h = \Phi_{h-1} \mathbb{A}_1 + \Phi_{h-2} \mathbb{A}_2$ for h-period ahead. Define a $k \times 1$ vector $\boldsymbol{e}_\lambda$, all of whose elements are zero except for a unit corresponding to the λ-shock. Given the impact matrix $\boldsymbol{B}^\star$, the impulse vector corresponding to the λ-shock is $a_\lambda = \boldsymbol{B}^\star \boldsymbol{e}_\lambda$ and the associated IRFs at horizon h are given by $\Gamma_{\lambda,h} = \Phi_h \boldsymbol{a}_\lambda$.*

2.2   The γ-Shock

To construct the γ-shock, I integrate the return-forecasting framework of the empirical finance literature (see Lettau and Ludvigson (2010); Cochrane (2011) for a review) with the VAR model 2.1. To estimate variation in the conditional mean of excess returns, the finance literature typically estimated the following predictive regression model:

$$r_{t+1}^H = a + \beta X_t^\star + \varepsilon_{t+1}, \tag{2.5}$$

where r_{t+1}^H is the cumulative log excess market return between $t+1$ and $t+H$; $X_t^\star$ is a vector of variables at the end of t used to predict the excess returns, and ε_{t+1} is a zero-mean disturbance term. Horse race among predictors is typically assessed using the R^2-statistic of the estimated regression 2.5. In my empirical model, $X_t^\star$ will be a subset of the state vector, X_t in 2.1. I partition the state vector as $X_t = \begin{bmatrix} \bar{X}_t; & X_t^\star \end{bmatrix}$ (in the spirit of Adrian, Crump, and Moench (2015)), where $\bar{X}_t$ are the remaining variables in the VAR that are not used in the predictive regression 2.5. Constructing the γ-shock is then based on the historical decomposition of $X_t^\star$ in the VAR.

Definition 6 (Historical Decomposition) *Consider the general VAR(p) model 2.1, and let $\upsilon_{i,t}(h) = X_{i,t+h} - E_{t-1} X_{i,t+h}$ be the h-step ahead forecast error in the i-th variable of the VAR. Given a structural impact matrix $\boldsymbol{B}$ and corresponding orthogonal shocks $f_t = \eta_t \boldsymbol{B}^{-1}$, the historical decomposition of $\upsilon_{i,t}(h)$ can be computed as the sum of the contribution of the k orthogonal shocks:*

$$\upsilon_{i,t}(h) = \sum_{j=1}^{k} \Gamma_{j,h} f_{j,t+h}, \tag{2.6}$$

where $\Gamma_{j,h}$ is the impulse response function corresponding to the j-th orthogonal shock in the VAR.

---

[9]For recent contributions, see Bryzgalova and Julliard (2015) and Boons and Tamoni (2015) amongst others.

The γ-shock is constructed to be an orthogonalised shock in the VAR which generates counterfactual variation $\hat{X}_t^\star$ in the predictor variables, such that when $\hat{X}_t^\star$ is used in the predictive regression 2.5, the associated R^2-statistic is maximised. In essence, this method finds the γ-shock by using the return-predictability step to restrict the historical decomposition of $X_t^\star$. I now briefly discuss how the γ-shock and the method in general can be connected to the macroeconometrics and finance literatures.

In terms of the link to the macroeconometrics literature, this strategy blends elements of recent developments on the identification of news shocks (Uhlig 2004; Barsky and Sims 2011; Kurmann and Otrok 2013). These papers identify news shocks about future economic fundamentals, based on maximising the contribution of these shocks to the forecast error variance of a selected variable in the VAR over a pre-specified future horizon. My orthogonlisation scheme is based on finding counterfactual variation in selected variables in the VAR that have maximal forecasting power as return predictors. In essence, the return-predictability step serves as an external reference point which restricts the entire sequence of historical decompositions in the VAR.

In terms of the link to the finance literature, the method aims to uncover the stochastic macroeconomic drivers of time-variation in risk premia. Since Fama and French (1989) and Ferson and Harvey (1991), growing empirical evidence points to the countercyclical nature of expected excess returns, implying that risk premia are high in recessions and low in expansions. While reduced-form macroeconomic variables such as the term spread (Fama and French, 1989) or the short-term interest rate (Fama and Schwert, 1977) have been found to forecast excess returns, it is clear that not all time-series variation in the term spread or the interest rate is driven by unexpected macroeconomic shocks that would ultimately lead to recessions and to spikes in risk premia. By constructing the γ-shock, the aim is to use a non-restrictive way to find the portion of variation in predictor variables that can be directly attributed to macroeconomic disturbances which cause recessions.

2.3   The Recession-Shock

As discussed above, finding the γ-shock which drives time-variation in expected returns may uncover the macroeconomic drivers of recessions. As a robustness exercise, I propose a third orthogonalisation scheme which directly looks for the drivers of recessions without using information on the time-variation in risk premia. The obtained Recession-shock is assumed to be the sole orthogonal macroeconomic contributor to a given historical event, e.g. to the Great Recession. One can then check how this macroeconomic force compares to the γ-shock, i.e. the driver of time-varying risk premia.

Methodologically, the identification of the Recession-shock is based on finding a rotation matrix Q such that one of the impulse vectors associated with the structural impact matrix $\boldsymbol{B}$ will deliver an orthogonal shock with the following property: the historical contribution

of this shock to a variable of interest in the VAR is as close as possible to the realised path of this variable over a given horizon. More formally, let Θ denote the set of possible rotations, Y_j is one of the variables in the VAR whose realised path the identified shock is to explain over a pre-specified period, with t_1 and t_2 denoting the start and end of the period. Let $\hat{Y}_{j,t_1:t_2}^{Recession-Shock}$ denote the counterfactual path of variable Y_j which would have realised between time t_1 and t_2, if the only source of business cycle fluctuations had been the Recession-shock. The algorithm is then written as:

$$Q_{opt} = \underset{Q \in \Theta}{\arg \min} \left| Y_{j,t_1:t_2} - \hat{Y}_{j,t_1:t_2}^{Recession-Shock}(Q) \right|. \tag{2.7}$$

In my application, Y_j will be aggregate consumption growth and the period of interest will be t_1=2008Q1, t_2=2009Q4. The problem 2.7 can be solved using numerical optimisation techniques. The method draws on the technique of matching impulse response functions (Christiano, Eichenbaum, and Evans, 2005), as well as on recent orthogonalisation schemes that use event-related restrictions (Ludvigson, Ma, and Ng, 2015, 2017; Antolin-Diaz and Rubio-Ramirez, 2018; Ben Zeev, 2018). The proposed orthogonalisation scheme is general, and could be used to explore whether the driving force of one particular historical event can account for the causes of other, seemingly similar, historical events as well.

3   Empirical Results

3.1   Data

The VAR includes quarterly data on output, consumption, price level, the short-term interest rate, the default spread, and the term spread. Consumption is total personal consumption expenditure from Greenwald, Lettau, and Ludvigson (2015). Output is seasonally adjusted real GDP (FRED code: GDPC1). Price level is the consumer price index for all urban consumers (FRED code: CPIAUCSL). Interest rate is the Federal Funds Rate (code: FED-FUNDS). Default spread is the difference between the AAA (FRED code: AAA) and BAA (FRED code: BAA) corporate bond yields. The term spread is the difference between the ten-year Treasury and T-bill rates as in Goyal and Welch (2008). The full sample period is 1963Q3-2015Q3, but I will also experiment with a shorter sample (1963Q3-2008Q3) that excludes the period of zero lower bound on the nominal interest rates. This brings my analysis closer to the information set that the monetary policy literature typically used to estimate policy shocks.

For test assets, I combine the standard 25 Size-B/M portfolios of Fama and French (1993) with the 30 industry portfolios. Returns are quarterly, in excess of the T-bill rate. Augmenting the FF25 with the 30 industry portfolios follows prescription 1 (pp. 182) of Lewellen, Nagel, and Shanken (2010), thereby relaxing the tight factor structure of Size-

B/M. For robustness, I will also use the 25 portfolios double sorted on size-profitability and size-investment, as these portfolios feature prominently in recent empirical asset pricing studies (Fama and French, 2015, 2016).

As return predictors, I use the short-term interest rate (Fama and Schwert, 1977), the term spread and the default spread (Fama and French, 1989); these macroeconomic variables featured prominently in predictive regressions. To construct the γ-shock, I use CRSP based S&P 500 return over the corresponding T-bill rate (from Goyal and Welch (2008)) as the regressand in the return predictability step ins 2.5. I will compare the return predictability performance (of the counterfactual time-series of the interest rate, the term spread and the default spread, implied by the γ-shock) to the model which uses the CAY variable proposed by Lettau and Ludvigson (2001a,b).[10]

The VAR is estimated in levels as most monetary policy VARs (Sims 1980; Christiano, Eichenbaum, and Evans 1999; Uhlig 2005; Gertler and Karadi 2015), thereby I avoid making any transformation to detrend the data (Sims, Stock, and Watson, 1990). My results change very little when I include a linear trend in the VAR. The model has two lags, based on the SIC. I will also cross-check that my results are robust to changing the lag length in the VAR.

3.2 The λ-Shock

The Aggregate Effects of the λ-Shock First, I present the results for the VAR estimates of the λ-shock, implied by standard equity portfolios. Using the OLS estimates, I compute IRFs for the λ-shock along with the IRFs associated with interest rate innovations using Cholesky-orthogonalisation.

The blue crossed lines in Figure 2 display the IRFs of the six-variables to a λ-shock that induces a 100bp increase in the interest rate.[11] The black circled lines correspond to a 100bp interest rate shock using Cholesky factorisation. This orthogonalisation method has been frequently used in the macroeconomics literature to identify monetary policy shocks (Sims 1980; Christiano, Eichenbaum, and Evans 1999). The lines in Figure 2 are difficult to tell apart, as the point estimates of the two sets of IRFs are virtually identical. The striking resemblance between the two sets of IRFs occurs in spite of the fact that constructing the λ-shock does not rely on any of the strong restrictions that Cholesky-identified monetary policy shocks traditionally relied on.[12]

The response of all variables is persistent. Note that the λ-shock has virtually no effect on consumption on impact, but the effect increases substantially with the horizon, reach-

---

[10]This variable measures deviations from a cointegrating relation for log consumption (C), log asset wealth (A) and log labour income (Y), and has proved to be successful in predicting excess stock market returns.

[11]As shown by Figure 9 in the Appendix, the shape of these IRFs is similar when the lag length is changed or when the Great Recession is excluded from the sample.

[12]The zero-restrictions under Cholesky-identification imply that the variables ordered before the monetary policy instrument do not respond to the monetary policy shock contemporaneously. See Section 4.1 of Christiano, Eichenbaum, and Evans (1999) for a detailed discussion of this recursiveness assumption.

ing a peak impact of about -0.5% approximately 12-15 quarters after the shock hits. This is consistent with the consumption dynamics implied by recent asset pricing models that highlighted the irrelevance of short-term consumption innovations to pricing (Bansal and Yaron, 2004; Parker and Julliard, 2005). The recent empirical evidence confirmed that slow moving consumption risk can explain the cross-sectional variation of average returns (Bryzgalova and Julliard, 2015; Boons and Tamoni, 2015). My framework adds to the literature by accounting for possible general equilibrium relationships between consumption growth and other macroeconomic variables while jointly explaining the cross-section of average returns. To that extent, the framework sheds further light on the macroeconomic drivers of the (reduced-form) slow consumption adjustment shocks, recently analysed by Bryzgalova and Julliard (2015), that can explain the cross-section of average returns.

Figure 2: Impulse Responses to a λ-shock and to a Cholesky Interest Rate Shock
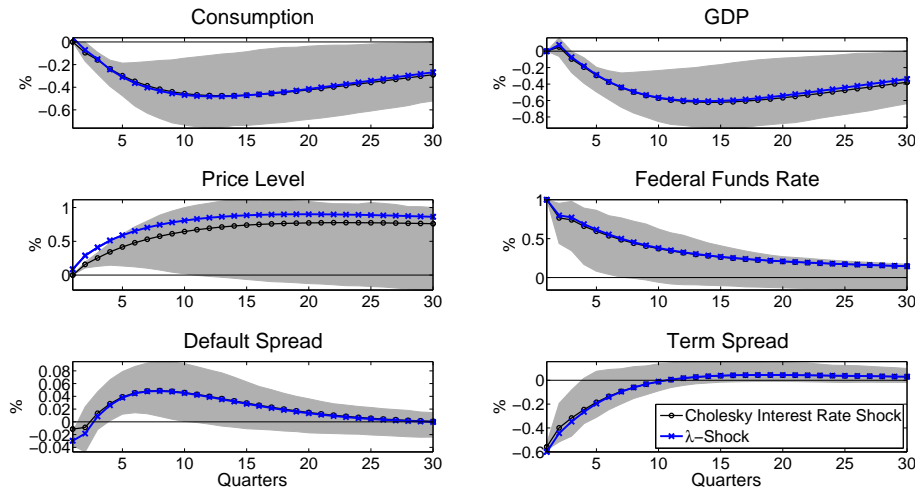

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The VAR(2) is estimated on a subsample 1963Q3-2008Q3. I use the FF55 to construct the λ-shock, whose IRFs are depicted by the black circles lines. The blue crossed lines are λ-shock, and the blacked circles lines are Cholesky-orthogonalised interest rate shock with the associated 95% confidence band (using wild-bootstrap). The IRFs are normalised to increase the interest rate by 100bp.

An other important feature of Figure 2, that extends results from reduced-form consumption based asset pricing models, is that the λ-shock generates a *negative comovement* between the nominal interest rate and consumption. In contrast, the recent macroeconomics literature on business cycles (Smets and Wouters, 2007; Christiano, Motto, and Rostagno, 2014; Negro, Eggertsson, Ferrero, and Kiyotaki, 2017) has found that the main drivers of economic downturns (including the Great Recession) generate a positive comovement between the nominal interest rate and consumption. Figure 2 is therefore important because it points to a possible dichotomy between the macroeconomic drivers of the cross-section of average returns and the drivers of recessions. This will be further investigated in the next Section.

Relation to Identified Monetary Policy Shocks   To highlight the relation of the λ-shock to the monetary policy literature, I compare the time-series of the λ-shock to other

benchmark estimates of monetary policy shocks. Figure 3 plots the time-series of the λ-shock against the policy shock series identified by Sims and Zha (2006). The correlation between the two series is 0.84. As a robustness exercise, I also check the correlation with narrative measures of monetary policy shocks: based on the overlapping estimation period 1969Q1–2007Q4, the correlation coefficient between the monetary policy shock series as identified in Romer and Romer (2004) [and updated by Tenreyro and Thwaites (2016)] and the λ-shock series is 0.75.

Figure 3: The λ-shock and Monetary Policy Shocks
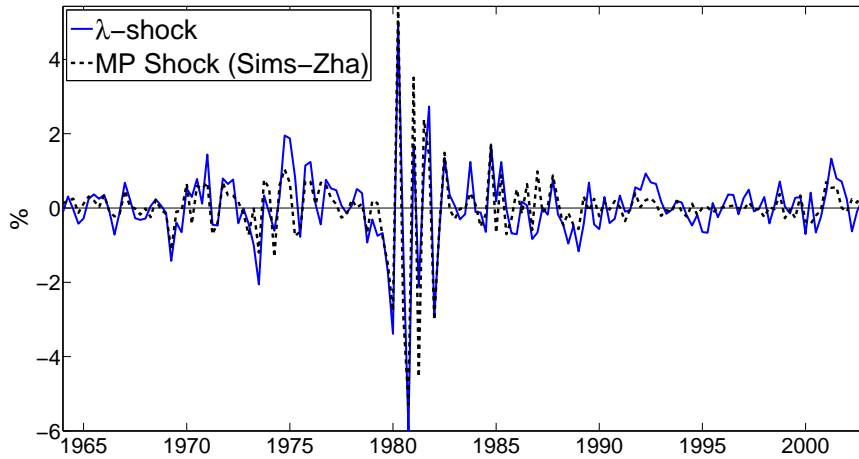
Notes: The λ-shock is from a six-variable VAR which includes quarterly data on consumption, GDP, CPI, Fed Funds rate, the term spread, and the default spread. The sample period is 1963Q3-2015Q3.Details of the orthogonalisation scheme are explained in the text below. The test assets are the 25 portfolios of Fama and French (1993) augmented with the 30 industry portfolios as prescribed by Lewellen, Nagel, and Shanken (2010). The monetary policy shock series are from Sims and Zha (2006), as documented in Stock and Watson (2012), and transformed to have unit standard deviation. The correlation coefficient is 0.84.

As an additional cross-check, I also employ the latest methodology of the monetary policy literature that uses high frequency asset price movements around policy announcements (Gurkaynak, Sack, and Swanson, 2005) as instruments for monetary policy shocks in a proxy-SVAR framework (Stock and Watson, 2012; Mertens and Ravn, 2013). To that end, I use the monthly VAR(12) model of Gertler and Karadi (2015) estimated over the 1979m7-2012m6 period. The correlation between the monetary policy shocks series from their four-variable VAR(12) model that generated their baseline Figure 1 (p. 61 of Gertler and Karadi (2015)) and the λ-shock implied by the FF55 is about 0.79.

Despite these findings, using off-the-shelf monetary policy shock series to price the cross-section would likely lead to the rejection of the corresponding pricing model, because of restrictive identifying assumptions and mis-measurement in macroeconomic data. This is an important point which provides justification for the somewhat reverse direction taken in this paper, whereby I start with asset prices and then work "backwards". This will be further discussed and illustrated by a Monte-Carlo exercise in Section 4 and in Section A.6 of the Appendix.

Relation to Other Identified Macroeconomic Shocks   The high empirical correlation between the λ-shock and VAR-based monetary policy shocks is non-trivial given that the number of other structural shocks that the literature has identified (such as technology shocks, tax shocks, government spending shocks amongst others, as recently reviewed by Ramey (2016)) are assumed to be orthogonal to policy shocks and, in theory, they all should affect the representative household's utility function and thereby correlate with SDF innovations. Based on my investigation of these shocks as collected by Ramey (2016), the estimated λ-shock has little empirical correlation with these shocks.[13]

However, a notable exception is news-type shocks to total factor productivity (TFP), as identified by Kurmann and Otrok (2013). Their Figure 4 shows IRFs for identified TFP news shocks that are similar to Figure 2. Based on my calculation for the overlapping sample (1963Q4–2005Q2), the correlation between their estimated TFP news shock series and the λ-shock series is 0.79. The ambiguity evoked by this result may seem an awkward outcome: after all, how can the λ-shock have such a high correlation with two, seemingly distinct structural disturbances? One simple explanation for such an ambiguity: TFP news shocks and monetary policy shocks are highly correlated in the data. Section B.10 of the Appendix provides strong evidence for this which, to the best of my knowledge, has not been documented in the literature yet. It is to note that my paper does not take a stand on the correct identification of monetary policy shocks or TFP news shocks, therefore this empirical conundrum is somewhat unrelated to the main focus of this paper. However, given the importance of this issue and its usefulness for providing macroeconomic interpretations of the λ-shock, Section B.10 of the Appendix comments on the possible drivers of this result. My analysis suggests that this empirical regularity is a possible symptom of an identification problem in the literature.

Overall, my results so far can be interpreted as supportive of the importance of monetary policy surprises in driving cross-section of average returns (Ozdagli and Velikov, 2016). Alternatively, a negative reading is that standard measures of monetary policy shocks are not very well identified: "In the absence of an empirically useful dynamic monetary theory, at least we can require the impulse-response functions to conform to qualitative theory such as Friedman (1968). Most VARs do not conform to this standard. Prices may go down, real interest rates go up, and output may be permanently affected by an expansionary shock" (Cochrane, 1994, p. 300).

Further Discussion   Two additional points are worth mentioning about the results so far. First, as mentioned in the Introduction, my method can separately control for the macroeconomic information set that I recover the estimated SDF innovations from, and for

---

[13]For example, the delayed expansion of aggregate quantities in Figure 2 makes the shock clearly distinct from unanticipated technology shocks that would have an immediate impact on consumption and output, as studied for long by the Real Business Cycle (RBC) literature.

the aggregate risks that the given test portfolios proxy. The λ-shock and the interest rate shock in Figure 2 are estimated on the same macroeconomic information set, which enables me to draw conclusions that Cholesky-identified interest rate shocks closely resemble SDF innovations in an unconditional linear asset pricing framework. In contrast, by simply comparing SDF innovations implied by the 3-factor model of Fama and French (1993) (or in fact any other empirical finance models of the SDF) to monetary policy shocks estimated by macroeconomists, it would be difficult to tell the reason for any possible lack of comovement between tho series. Identified monetary policy shocks and SDF innovations implied by the 3-factor model may have very little comovement simply because the 3-factor model does not span the macroeconomic information space that would be an important input for the monetary policy reaction function and thereby for the identification of the non-systematic, surprise element of monetary policy.[14]

Second, related to the last point, the price level response in Figure 2 is counter to how monetary policy in New Keynesian models tend to affect prices (Gali, 2008; Woodford, 2003). This is the well-known 'price puzzle' associated with Cholesky orthogonalisation in VAR models (Sims, 1992).[15] My paper does not take a stand on the right identification of monetary policy shocks; it only shows that monetary policy shocks, as typically identified by macroeconomists, resemble the orthogonal shock which explains the cross section of expected excess returns. Nevertheless, I further investigate the issues related to the price response in Section B.2 of the Appendix, by re-estimating the λ-shock in the context of Uhlig (2005).[16]

3.3   The γ-Shock

I will now analyse the economic properties of the γ-shock. Recall that this is an orthogonalised shock in the VAR which is constructed to explain time-variation in risk premia. The key step in this methodology is the nested estimation of the return forecasting regression (2.5). Given their popularity as return predictors since at least Fama and French (1989), I will employ the last three variables in the VAR (FFR, default spread and term spread) as predictors of the conditional mean of excess stock market returns.

Forecasting Excess Returns   Before analysing the dynamic effects of the γ-shock on the macroeconomy, I first discuss the results of the predictive regression step, summarised in Table 1. I construct the γ-shock by maximising the corresponding return forecasting power

---

[14]As will be discussed in Section 4, the pricing performance of the λ-shock implied by my baseline model is approximately on par with the pricing performance of the 3-factor model; and to increase the pricing performance of the λ-shock, one can simply augment the VAR with standard pricing factors.

[15]See Ramey (2016) for a recent discussion.

[16] Uhlig (2005) uses a monthly VAR model and imposes sign restrictions on the impulse response functions to identify monetary policy shocks. This identification scheme, by construction, "solves" the price-puzzle. While the empirical model of Uhlig (2005) is different from my baseline (different sample period, monthly instead of quarterly frequency), I continue to find that λ-shocks resemble monetary policy shocks (see Figure 11 of Appendix).

at 4-quarter horizon, and (using the same γ-shock) I compute the results for different horizons ranging from one quarter ahead up to two years ahead. Table 1 reports the regression coefficients, t-statistics using the Hansen and Hodrick (1980)-correction, and the adjusted R^2 statistics for each regression. Panel A reports the results using the actual VAR variables as predictors; Panel B shows the results using the CAY variable of Lettau and Ludvigson (2001a,b); Panel C reports the results using the three counterfactual VAR variables induced by the γ-shock; Panel D reports the results using the counterfactual variables induced by all other shocks that are orthogonal to the γ-shock.

Panel A and Panel B of Table 1 confirm the evidence of the macro-finance literature (reviewed by Lettau and Ludvigson (2010) recently) on the superiority of CAY, as a return predictor, over the short-term interest rate, the default spread and the term spread. For example, the last column of the table shows that the CAY variable explains around 23% of two-year ahead excess stock market returns. In contrast, the regression that includes the last three variables of my baseline VAR only explains 11% of excess returns at the same horizon.

However, as discussed above, not all variation in the these macroeconomic variables are related to future recessions; and using all the variation in these reduced-form variables may therefore not predict future returns very well. Panel C of Table 1 confirms that using the counterfactual time-series, induced by the γ-shock, as predictors, substantially improves the forecasting power of the predictive regression, explaining around 40% of two-year ahead excess returns. At almost every horizon, the variation in the interest rate, default spread and the term spread, induced by the γ-shock, explains more than twice as much of future excess returns as the regression model which uses CAY. Panel D shows the results when all the remaining variation in the three macroeconomic variables (not explained by the γ-shock) is used in predicting future excess returns. As expected, this variation is not useful in predicting returns with all adjusted R^2 statistics being around zero at all forecast horizons.[17]

To further illustrate the reduced-form nature of return predictors, Figure 13 in Appendix B.4 plots the time-series of the term spread along with the counterfactual time-series of the term spread induced by the γ-shock – the variation most relevant to predicting excess returns. While the two series are correlated (0.75), the correlation clearly breaks down in certain periods such as the 1980s. In fact, the historical decomposition will show that aggregate fluctuations during this period were mainly driven by the other orthogonal force in the model, the λ-shock. Overall, Figure 13 illustrates that, while the term spread "tends to be low near business-cycle peaks and high near troughs" (Fama and French, 1989), not all business cycles (and variation in return predictors) have been caused by the macroeconomic

---

[17] Similarly, one could apply this type of analysis to investigating the macroeconomic shocks underlying time-varying bond premia. My investigation of the dataset of Ludvigson and Ng (2009) suggests that using the counterfactual variation in their principal components, induced by the corresponding γ-shock, increases their R^2 statistics by about 30%; and the λ-shock continues to generate a sharp response in their first principal component (related to real activity). These results are available on request.

	1Q	2Q	3Q	4Q	5Q	6Q	7Q	8Q
				Forecast Horizon H				
	colspan							

	1Q	2Q	3Q	4Q	5Q	6Q	7Q	8Q
Model A: Actual VAR Variables								
FFR	-0.22	-0.32	-0.27	-0.22	-0.07	0.20	0.49	0.74
	(-0.86)	(-0.76)	(-0.50)	(-0.36)	(-0.10)	(0.25)	(0.51)	(0.68)
DEF	1.53	3.46	3.92	4.15	3.80	3.44	2.75	1.83
	(0.77)	(1.18)	(1.13)	(1.10)	(0.92)	(0.72)	(0.45)	(0.25)
TERM	0.19	0.54	1.36	2.11	2.90	4.05	4.98	5.95
	(0.37)	(0.60)	(1.10)	(1.38)	(1.70)	(2.22)	(2.51)	(2.85)
	[0.01]	**[0.02]**	**[0.04]**	**[0.06]**	**[0.07]**	**[0.08]**	**[0.09]**	**[0.11]**
Model B: CAY								
CAY	0.55	1.12	1.77	2.52	3.17	4.00	4.92	5.77
	(2.18)	(2.25)	(2.33)	(2.46)	(2.59)	(3.00)	(3.58)	(4.25)
	[0.01]	**[0.03]**	**[0.05]**	**[0.09]**	**[0.11]**	**[0.14]**	**[0.19]**	**[0.23]**
Model C: Counterfactual VAR Variables Induced by the γ-Shock								
FFR	0.59	1.14	1.66	2.15	2.61	3.15	3.66	4.20
	(1.21)	(1.14)	(1.19)	(1.25)	(1.32)	(1.44)	(1.52)	(1.66)
DEF	-9.14	-16.87	-23.42	-29.06	-33.96	-38.78	-43.34	-48.31
	(-2.01)	(-1.73)	(-1.78)	(-1.86)	(-1.92)	(-2.04)	(-2.15)	(-2.33)
TERM	3.22	6.47	9.63	12.66	15.46	18.41	21.07	23.65
	(2.53)	(2.42)	(2.42)	(2.46)	(2.53)	(2.63)	(2.73)	(2.93)
	[0.04]	**[0.08]**	**[0.14]**	**[0.20]**	**[0.25]**	**[0.30]**	**[0.34]**	**[0.39]**
Model D: Counterfactual VAR Variables Induced by All Other Shocks								
FFR	-0.16	-0.21	-0.13	-0.06	0.10	0.36	0.62	0.84
	(-0.64)	(-0.47)	(-0.21)	(-0.08)	(0.11)	(0.33)	(0.47)	(0.55)
DEF	1.76	3.67	3.80	3.49	2.52	1.63	0.53	-0.64
	(0.88)	(1.20)	(0.93)	(0.73)	(0.48)	(0.28)	(0.07)	(-0.07)
TERM	-0.18	-0.43	-0.27	-0.23	-0.16	0.25	0.51	0.97
	(-0.30)	(-0.39)	(-0.17)	(-0.12)	(-0.07)	(0.10)	(0.18)	(0.31)
	[-0.01]	**[-0.00]**	**[-0.00]**	**[-0.01]**	**[-0.01]**	**[-0.01]**	**[-0.01]**	**[-0.00]**

Notes: The table reports results from regressions of excess returns on lagged variables. H denotes the return horizon in quarters. The dependent variable is the sum of H log excess returns on the CRSP based S&P Composite Index. The regressors are one-period lagged values of actual time-series of the federal funds rate (FFR), the term-spread (TERM) and the default spread (DEF) in Model A, the CAY measure of Lettau and Ludvigson (2001a,b) in Model B, and the counterfactual time-series of FFR, TERM and DEF (induced by the γ-shock) from a six-variable VAR(2) estimated over 1963Q3-2015Q3. The γ-shock is constructed so that the corresponding forecast power at the four-quarter horizon is maximised. For each of the three regressions, the table reports the OLS estimates of the regressors, the t-statistics using the Hansen and Hodrick (1980) correction (as implemented in Cochrane (2011)) are in parentheses, and adjusted R^2 statistics are in the bolded square brackets. Both the CAY measure and the counterfactual predictors are treated as known variables.

force which drives time-variation in expected returns.

Naturally, the same logic applies to the estimated time-series of expected excess returns implied by the given predictors. The lower panel of Figure 13 shows the time-series of realised cumulative 8-quarter excess returns (dashed line) against the return forecast implied by the actual time-series of the predictors (dotted line) and the forecast implied by the counterfactual time-series of the predictors induced by the γ-shock (solid line). The coefficients of correlation between realised returns and the data-based forecast and the counterfactual-

based forecast are 0.36 and 0.63, respectively. This is another way of conveying the message summarised in Table 1: variations in return predictors are driven by a range of macroeconomic forces, and not all these forces change the conditional mean of excess returns.

The Aggregate Effects of the γ-Shock and Relation to the Macroeconomic Literature To analyse the macroeconomic properties of the γ-shock, I compute impulse responses. Recall that, to the extent that time-variation in expected returns is linked to economic booms and busts (Lettau and Ludvigson, 2010; Cochrane, 2011), the γ-shock can be thought of as the stochastic driver of recessions in the VAR. Figure 4 shows the impulse responses for the γ-shock along with the responses for the λ-shock. Both shocks are set to be contractionary in the Figure. The behaviour of the γ-shock is distinct: it causes a sharp drop in consumption and output, and generates a positive comovement among quantities and short-term interest rate. The fall in the short-term rate is indicative of the monetary policy authority endogenously responding to the recessionary γ-shock by loosening policy (Taylor, 1993). Moreover, the default spread and the term spread immediately jump up in response to a γ-shock, whereas the λ-shock causes a delayed increase in both quantity variables and in the default spread, and generates a fall (rather than a rise) in the term spread. Overall, the dynamics generated by the γ-shock resemble features of recent US recessions.

Figure 4: Impulse Responses to a γ-Shock and to a λ-Shock
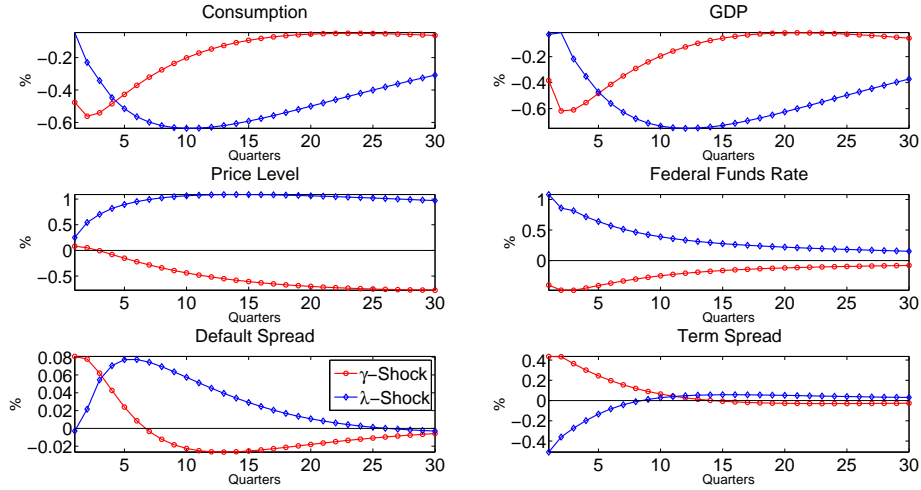
Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The VAR(2) is estimated on the sample 1963Q3-2015Q3. I use the FF55 to construct the λ-shock, and use the CRSP based aggregate SP500 stock market return (over the corresponding T-bill rate) to construct the γ-shock.

I now briefly discuss how the γ-shock is related to recent macroeconomic explanations of the business cycle. This literature has long sought to construct general equilibrium models that could generate the type of comovements that are induced by the γ-shock, as in Figure 4. The seminal macroeconomic paper by Smets and Wouters (2007) built a New Keynesian general equilibrium model which was among the first that could explain these empirical

features of the data. In this model, a large fraction of short-term consumption fluctuation is driven by demand-type shocks including disturbances ("preference shocks") that directly distort the representative household's marginal utility (analogous to Y_{t+1} in equation 1.1).[18]

The most recent macroeconomic papers such as Christiano, Motto, and Rostagno (2014) combined standard New Keynesian features with a model of financial intermediation, and also used financial data (in addition to macroeconomic data) for the model estimation. This paper shows that fluctuations in the severity of Bernanke, Gertler, and Gilchrist (1999)-type agency problems associated with financial intermediation explain up to 60% of US business cycle fluctuations. While the model of Christiano, Motto, and Rostagno (2014) is linearised and thereby absent of time-varying risk premia, the driving force in their model is exogenously fluctuating uncertainty related to the cross-section of idiosyncratic production risk, which they refer to as "risk shocks". Their results suggest that risk shocks are the primitive macroeconomic force that is proxied by preference shocks in models without a financial sector.

Figure 5: γ-Shocks and Risk-shocks of Christiano, Motto, and Rostagno (2014)
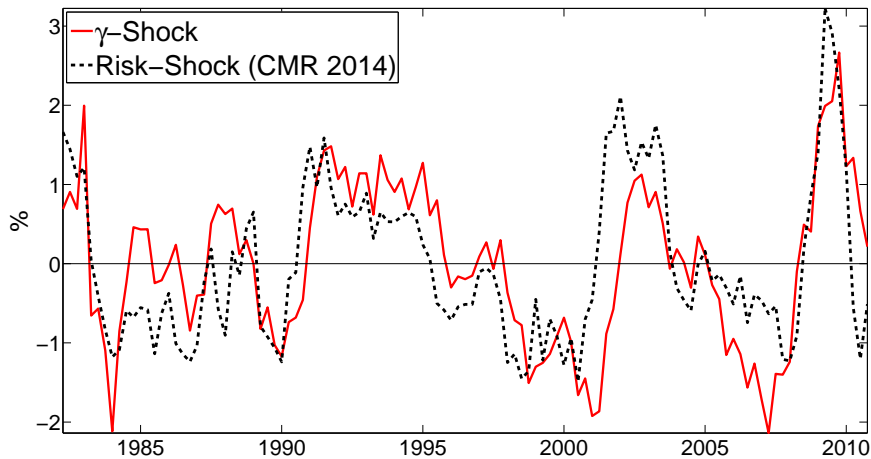
Notes: The figure shows the four-quarter moving average of the risk shock of CMR (2014) and the (inverse of the) γ-shock, estimated from my baseline six-variable VAR over 1963Q2-2015Q3, using the CRPS based aggregate SP500 stock market return (over the corresponding T-bill rate). The correlation coefficient is 0.67.

Given the ability of risk shocks to explain most recent US recessions, I compare the time-series of the risk shock series of Christiano, Motto, and Rostagno (2014) to the estimated time-series of the γ-shock from my VAR model. Figure 5 shows that estimated time-series of risk shocks bear a close resemblance with the γ-shock. This is despite the fact that the two shocks series are estimated using different information sets and very different methodologies. Both shocks made a sharp contribution to the recessions in the early 1990s, early 2000s and

---

[18]The role of demand-type preference shocks (i.e. innovations in the state variable Y_{t+1} in equation 1.1) in driving the business cycle is also important in other New Keynesian models (without a financial sector and financial shocks) such as Christiano, Eichenbaum, and Evans (2005). The estimates shown in Table 5 of Christiano, Motto, and Rostagno (2014) suggest that these shocks explain about 67% of aggregate consumption at business cycle frequency.

in the Great Recession. These results suggest that risk shocks are not only important in contributing to macroeconomic fluctuations but also drive time-variation in aggregate risk premia. At a more general level, these results highlight that integrating the return forecasting framework of the empirical finance literature (Campbell and Shiller 1988; Fama and French 1989; Goyal and Welch 2008; Pastor and Stambaugh 2009; Lettau and Ludvigson 2010; Cochrane 2011) with the structural VAR methodology (Sims, 1980) via the γ-shock can be useful for business cycle analysis.

The Orthogonality of the γ-Shock to the λ-Shock   So far, I have analysed the λ-shock and the γ-shock separately, without making assumptions about the covariance structure of these two shocks. The corresponding IRFs (Figure 4) suggest that they capture different macroeconomic forces. To formally check the possible orthogonality of these two shocks with respect to one another, I compare the IRFs obtained by implementing the orthogonalisation schemes separately to the IRFs obtained by implementing the orthogonalisation schemes jointly. Figure 14 in the Appendix shows that both sets of IRFs are virtually identical, confirming that the two shocks can be regarded as orthogonal to each other.

This is an important empirical result because it shows that the macroeconomic shocks that determine the *level* of risk premia are distinct from the shocks that determine *time-variation* in risk premia. Moreover, the orthogonality of the λ-shock and the γ-shock enables the use of structural VAR decompositions, which can help quantify the historical contribution of these shocks to business cycle fluctuations.

3.4   Explaining the Business Cycle

To assess the contribution of the λ-shock and the γ-shock to US business cycles, I compute FEV decompositions over different forecast horizons. Table 2 shows that the λ-shock explains less than 1% of consumption fluctuations over the one-quarter horizon, but the contribution rapidly increases with the horizon and the shock explains around 60-65% of fluctuations over the longer (4-8 years) horizon. The λ-shock explains more than 70% of interest rate fluctuations over the 1-4 quarter horizon and the contribution falls only little at longer-frequency. This provides additional illustration that λ-shocks are empirically related to interest rate innovations studied by macroeconomists (Sims 1980; Christiano, Eichenbaum, and Evans 1999; Ramey 2016).

Table 2: The Contribution of the λ-shock and γ-shock to Business Cycles: FEV Decomposition

	Consumption				Federal Funds Rate			
	λ-Shock	γ-Shock	$\lambda\&\gamma$	Unexpl.	λ-Shock	γ-Shock	$\lambda\&\gamma$	Unexpl.
1Q	0.3	67.4	67.7	32.3	77.6	11.6	89.2	10.8
2Q	6.0	66.4	72.3	27.7	73.4	16.1	89.6	10.4
3Q	11.9	61.6	73.5	26.5	71.0	18.3	89.4	10.6
4Q	18.7	55.6	74.3	25.7	69.4	19.6	89.0	11.0
8Q	40.8	36.8	77.6	22.4	66.4	21.6	88.0	12.0
16Q	60.5	21.2	81.7	18.3	64.9	22.0	86.9	13.1
32Q	66.2	14.8	81.0	19.0	64.0	21.6	85.7	14.3

Notes: The table shows the % fraction of the total forecast error variance that is explained by the λ-shock and the γ-shock over different forecast horizons. The FF55 portfolios are used as test portfolios for the VAR model.
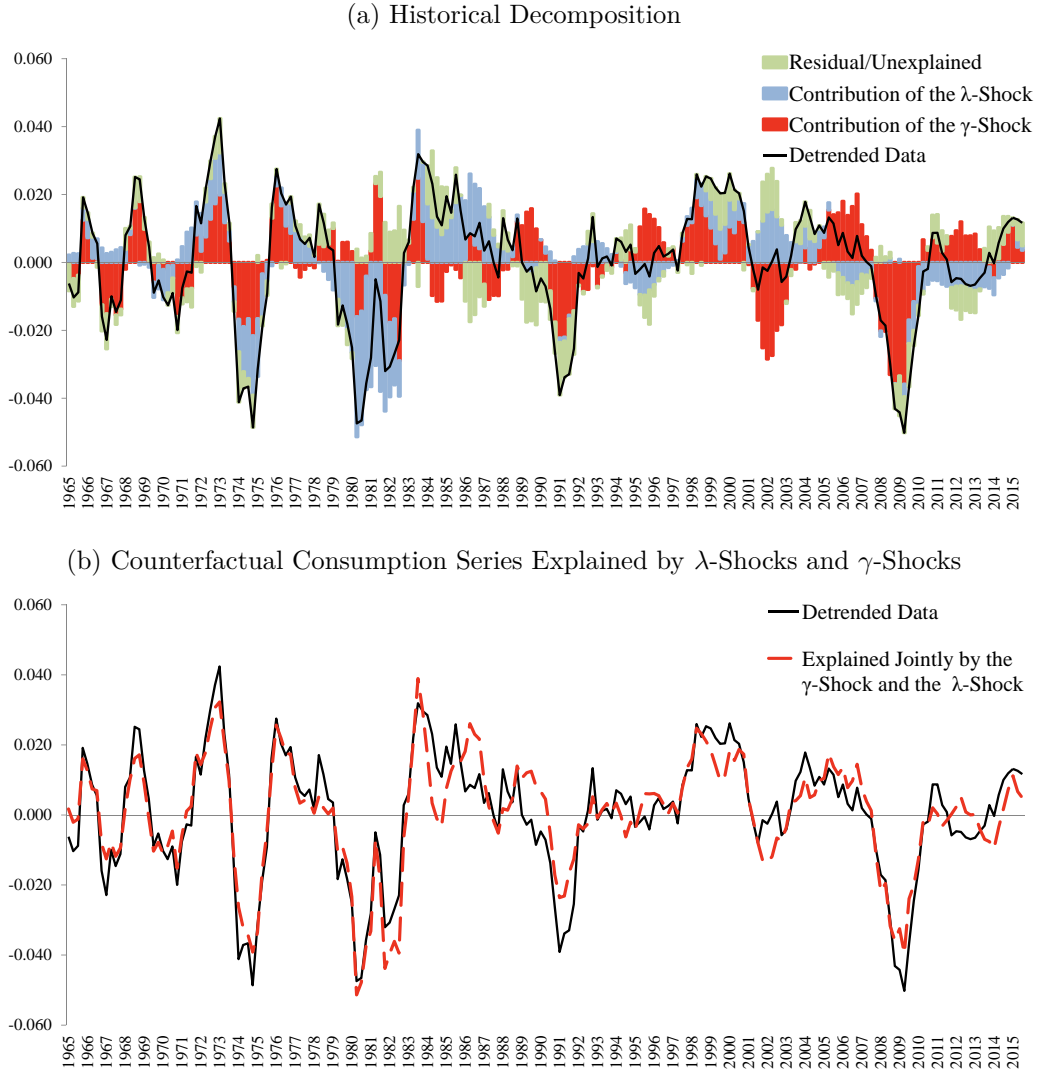
In contrast, the γ-shock explains a large fraction of short-term fluctuations in consumption, and only moderately contributes to the forecast error variance in the short-term interest rate. Overall, the joint contribution of the λ-shock and the γ-shock to consumption and interest rate fluctuations amounts to 70-90% at business cycle frequency.

Another way to assess the importance of these two shocks to business cycle fluctuations is to compute historical decompositions. This is shown in Figure 6. The black solid line in Panel A shows year-on-year consumption growth after removing the deterministic trend implied by the VAR. The contribution of the λ-shocks and the γ-shocks is represented by the blue and red bars, respectively; whereas the green bars show the contribution of the remaining residual disturbances in the VAR. The results show that the λ-shock contributed largely to the recession in the early 1980s, and to a smaller extent to the recession in 1974-75. All other downturns including the Great Recession can be explained by the γ-shock. Moreover, consistent with the FEV results, Panel B of Figure 6 shows that most historical fluctuations in consumption growth, over the past 50 years, can be jointly explained by the λ-shock and the γ-shock.

The fact that merely two orthogonal macroeconomic forces can explain the bulk of aggregate consumption fluctuations is a notable result. However, in quantitative models of the business cycle it is not atypical to have two dominant shocks explaining such a large fraction of fluctuations. This is true for more atheoretical, VAR models such as Blanchard and Quah (1989) or highly structural models such as Christiano, Motto, and Rostagno (2014). What is more important about my results is that these two dominant shocks are constructed based exclusively on asset price information and not on macroeconomic assumptions.

At a deeper level, these results highlight the importance of asset pricing explorations for macroeconomics (Cochrane and Hansen, 1992). Modern macroeconomic models have mainly focused on understanding the dynamics of aggregate quantities, and information on the level of and variation in expected returns has been largely ignored. Since Mehra and Prescott (1985), these models have struggled to explain the level of and variation in risk

Figure 6: Decomposing Year-on-year US Consumption Growth: the Role of λ-Shocks and γ-Shocks

(a) Historical Decomposition

(b) Counterfactual Consumption Series Explained by λ-Shocks and γ-Shocks

Notes: The figure shows the results from the historical decomposition implied by the baseline six-variable VAR(2) estimated on the sample 1963Q3-2015Q3. I use the FF55 to construct the λ-shock, and the CRSP based SP500 return (in excess of the corresponding T-bill rate) to construct the γ-shock. The contribution of the other four unrestricted orthogonalised shocks is depicted as residuals. The deterministic trend component, implied by the VAR, is removed from the time-series.

premia.[19] Via its more atheoretical nature, the macroeconometric framework proposed in this paper is able to cut this Gordian knot. It continues to capture the rich dynamics of macroeconomic time-series (Sims, 1980) while connecting it with the study of asset prices.

I carry out a number of robustness checks. First, I check how the FEV decomposition changes as I increase the lag length of the VAR from two to three and four. Table 5 in the Appendix confirms that my results are robust to these perturbations. I also recompute the historical decomposition implied by a VAR(4) model (Figure 10) and find little difference in the model's interpretation of history. Moreover, I explore the contribution of the λ-shock and the γ-shock to variation in aggregate consumption at a lower frequency. In addition to

---

[19]As Cochrane (2011) explains: "The job is just hard. Macroeconomic models are technically complicated. Macroeconomic models with time-varying risk premia are even harder" (p. 1090).

decomposing year-on-year consumption growth (as in Figure 6), I also compute the historical decomposition of the deviation of the level of aggregate consumption from the trend implied by my baseline VAR model. Figure 12 in Appendix B.3 shows that, consistent with the FEV results (Table 2), the λ-shock explains more of the low-frequency variation in consumption, including the persistent expansion above trend in the run-up to the Great Recession.

4 Robustness and Extensions

I carry out a number of additional robustness checks, summarised by paragraphs in the main text and explained in detail in Appendix B.

The Recession-Shock and the γ-Shock   To further explore the relation between recessions and time-varying risk premia, I study the effects of the Recession-shock and compare them to those implied by the γ-shock. Recall that the Recession-shock is engineered to explain the consumption collapse during the Great Recession (without any reference to asset prices). This identification theme provides an agnostic way (i) to explore the dynamic effects of the shock that triggered the Great Recession without directly restricting the impact of the shock on macroeconomic variables, and (ii) to check whether previous recessions had been caused by the same macroeconomic force that triggered the Great Recession. It remains an empirical question whether the Recession-shock behaves similarly to the γ-shock. My findings are reported in Section B.6. Both the IRF analysis (Figure 15) and the historical decomposition (Figure 16) confirm that the two shocks proxy the same macroeconomic force.

Other Equity Portfolios and Government Bond Returns   I check how the behaviour of the λ-shock changes when the same VAR model and the orthogonalisation method are applied to other test assets. A natural choice is the 25 portfolios double sorted on size-profitability and size-investment. These portfolios feature prominently in the most recent empirical asset pricing studies (Fama and French, 2015, 2016). In addition, I compute the IRFs for the λ-shock implied by the benchmark FF25 portfolios, sorted on size-B/M, that have been the most studied test assets to date.[20] Moreover, I also study the λ-shock implied by the excess returns on US government bonds. Section B.7 of the Appendix confirms that the results are similar to the baseline. This is consistent with the relatively small but growing literature on the joint pricing of stocks and bonds (Lettau and Wachter 2011; Bryzgalova and Julliard 2015; Koijen, Lustig, and Van Nieuwerburgh 2017).

Changing the VAR: an Example via Financial Intermediary Balance Sheets   One can also change the state-variables in the VAR which can serve two purposes. First, given

---

[20]An earlier version of this paper (Pinter, 2016) contains further analysis of the λ-shock implied by momentum returns (Asness, 1994; Jegadeesh and Titman, 1993).

the proliferation of asset pricing factors in the finance literature (Harvey, Liu, and Zhu, 2016), the proposed VAR framework could model the joint dynamics of any reduced-form variables that individually have been found to price the cross-section of returns, and to link the common stochastic driver of these variables to a single or multiple orthogonal shocks. Second, one can explore how the realisation of aggregate risks proxied by the given test assets may affect different parts of the macroeconomy. To support both these points, I will add to the VAR the aggregate capital ratio of the financial intermediary sector (constructed by He, Kelly, and Manela (2017)). The health of financial intermediary balance sheets has be a focus of attention in the asset pricing (Adrian, Etula, and Muir, 2014) and the macroeconomics literatures (Gertler and Kiyotaki 2010; Brunnermeier and Sannikov 2014; He and Krishnamurthy 2014). Section B.8 of the Appendix analyses the effects of the λ-shock and the γ-shock on the intermediary capital ratio. This aims to illustrate how the present framework could increase our understanding of the empirical linkages between macroeconomic fluctuations and any reduced-form asset pricing factor proposed by the empirical finance literature.

Pricing the Cross-section of Stock Returns It is worth noting that the focus of this paper is not the asset pricing performance of the λ-shock. Conditional on the proposed VAR model being an accurate representation of the economy, λ-shocks are SDF innovations by construction. Put it differently, the pricing performance of the given λ-shock can easily be improved by augmenting the VAR (e.g. including standard pricing factors such as the excess return on the market[21]) but not by changing the orthogonalisation assumption. Though the baseline six-variable VAR is far from being an accurate representation of the economy, it is a standard and parsimonious way of summarising macroeconomic dynamics. For the interested reader, I do summarise in Section B.9 of the Appendix the asset-pricing performance of the λ-shock implied by each test portfolios studied above. Overall, the pricing performance of the VAR (or equivalently, the λ-shock) is comparable with the 3-factor model of Fama and French (1993).

Using Identified Macroeconomic Shocks Directly for Pricing A further question relates to the use of established external shock series (identified by the macroeconomics literature) to price the cross-section, instead of using the cross-section of asset prices to back out the λ-shock and comparing it to external macroeconomic shocks. The reverse direction, taken in this paper, is motivated by the fact that identification of macroeconomic shocks may suffer from overly restrictive identifying assumptions and from mis-measurement of macroeconomic data. This has particularly relevant asset pricing implications, given that different ways of identifying the same macroeconomic shock can lead to different estimated time-series of the given structural shock (thereby leading to hugely different pricing

---

[21]These results are available upon request.

performance) in spite of the fact that the given identification schemes may lead to similar impulse response functions, as discussed in the debate between Rudebusch (1998) and Sims (1998). These problems make it likely for the given macroeconomic shock to be rejected as a pricing factor, even though the shock may be truly correlated with SDF innovations. To illustrate this point, Section A.6 provides a Monte-Carlo simulation whereby I first construct the SDF from the space of excess returns, and then introduce and gradually increase the measurement error in the SDF. I show that mis-measurement in the candidate factor leads to a more rapid deterioration (and increasing uncertainty) in pricing performance than in the reduction of its correlation with the true SDF.

Results from the UK   To check whether the results are robust to countries other than the US, I re-estimate my empirical model using data for the UK by making use of the availability of comparable test assets (Dimson, Nagel, and Quigley, 2003) across the two countries, and building on recent empirical work on return predictability in the UK (Chin and Polk, 2015). Section B.11 of the Appendix shows the results for the λ-shock and the γ-shock implied by the UK data. The results are similar to those found for the US. Moreover, the estimated time-series of the λ-shock continuous to be empirically related to monetary policy shocks, as constructed by Cloyne and Hurtgen (2016) for the UK.

Relation to Asset Pricing Theories   An important empirical finding of my paper is that the λ-shock and the γ-shock seem orthogonal to each other in the data. While the λ-shock explains the average level of risk premia, the γ-shock drives variation in it. These results can be connected to consumption based asset pricing theories via the generalised discount factor (1.1). For example, in models of long-run risk (Bansal and Yaron, 2004), the main drivers of the level of risk premia are shocks to the long-run component of consumption dynamics. As discussed recently by Cochrane (2017), these shocks cannot produce time-varying risk premia, and shocks to economic uncertainty are needed in these models to generate time-variation in expected returns. My results can also be interpreted through other asset pricing theories as well. As argued by Cochrane (2017), "Many of the models also capture the same intuitions (...) Each of the models suggests different candidates for the state variable Y_{t+1}. But these candidates are highly correlated with each other, and each sensibly indicative of fear or bad economic times. Telling them apart empirically is not easy, and possibly not that productive." (p. 982) For example, Y_{t+1} may proxy time-varying cross-sectional variance of individual consumption growth (Constantinides and Duffie, 1996). This interpretation of the γ-shock also has advantages, as it may explain the close empirical relationship with risk shocks of Christiano, Motto, and Rostagno (2014) which is directly related to the time-varying volatility of cross-sectional idiosyncratic uncertainty. Overall, a key advantage of my empirical methodology is its reverse engineering nature: it uses asset prices to construct the drivers of business cycles, instead of first having to take a stand on the theoretical framework

before testing the implied asset price predictions.

5   Summary and Conclusion

To conclude, my paper makes economic as well as methodological contributions that can help narrow the gap between finance and macroeconomics, and also help answer two main questions of macro-finance, as posed by Cochrane (2017) in the opening quote.

The first economic contribution is to highlight the overlap between unconditional linear pricing models proposed by the empirical finance literature (Fama and French, 1993) and structural shocks identified by the macroeconometric literature (Sims, 1980). I find that the macroeconomic force that drives the cross-sectional variation in expected returns closely resembles VAR-based estimates of monetary policy shocks, as identified by the macroeconomics literature. A positive reading of this finding is that it provides further evidence on the importance of monetary policy surprises in driving risk premia (Ozdagli and Velikov, 2016). A negative reading of this finding is that standard measures of monetary policy shocks may not be well identified.

The second economic contribution of this paper is to show that time-variation in expected returns can be explained by a demand-type macroeconomic shock that can account for most recent US recessions, and which closely resembles risk shocks, identified recently by Christiano, Motto, and Rostagno (2014). Moreover, I provided empirical evidence that the macroeconomic drivers of the level risk premia are orthogonal to the drivers behind the time-variation in risk premia, and argued that this is consistent with consumption-based asset pricing models. While the construction of the λ-shock and the γ-shock relies exclusively on asset price information and not on macroeconomic assumptions, I find that these two shocks jointly explain up to 80% of aggregate consumption fluctuations in the US.

The main methodological contribution of this paper is to propose a unified empirical framework whereby structural VAR methods can be combined with asset price information on the cross-sectional and time-series variation in expected asset returns, in order to study the stochastic drivers of business cycle fluctuations and their links with the dynamics of risk premia. A smaller methodological contribution is the proposed method to identify Recession-shocks in a VAR, which allows one to explore whether the driving force of one particular recession can account for the causes of other recessions as well. Applying these methods to data from other markets and other countries would be an interesting avenue for future research.

References

ADRIAN, T., R. K. CRUMP, AND E. MOENCH (2015): "Regression-based estimation of dynamic asset pricing models," *Journal of Financial Economics*, 118(2), 211–244.

ADRIAN, T., E. ETULA, AND T. MUIR (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *The Journal of Finance*, 69(6), 2557–2596.

ANTOLIN-DIAZ, J., AND J. F. RUBIO-RAMIREZ (2018): "Narrative Sign Restrictions for SVARs," *American Economic Review*, Forthcoming.

ARIAS, J. E., J. F. RUBIO-RAMIREZ, AND D. F. WAGGONER (2018): "Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 86(2), 685–720.

ASNESS, C. (1994): "The power of past stock returns to explain future stock returns," Working paper, University of Chicago.

BANSAL, R., AND A. YARON (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59(4), 1481–1509.

BARSKY, R. B., AND E. R. SIMS (2011): "News shocks and business cycles," *Journal of Monetary Economics*, 58(3), 273–289.

BAUMEISTER, C., AND J. D. HAMILTON (2015): "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," *Econometrica*, 83(5), 1963–1999.

BEN ZEEV, N. (2018): "What can we learn about news shocks from the late 1990s and early 2000s boom-bust period?," *Journal of Economic Dynamics and Control*, 87(C), 94–105.

BERNANKE, B. S., M. GERTLER, AND S. GILCHRIST (1999): "The financial accelerator in a quantitative business cycle framework," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, chap. 21, pp. 1341–1393. Elsevier.

BERNANKE, B. S., AND K. N. KUTTNER (2005): "What Explains the Stock Market's Reaction to Federal Reserve Policy?," *Journal of Finance*, 60(3), 1221–1257.

BLANCHARD, O. J., AND D. QUAH (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79(4), 655–73.

BOONS, M. (2016): "State variables, macroeconomic activity, and the cross section of individual stocks," *Journal of Financial Economics*, 119(3), 489–511.

BOONS, M., AND A. TAMONI (2015): "Horizon-Specific Macroeconomic Risks and the Cross-Section of Expected Returns," mimeo, LSE.

BRUNNERMEIER, M. K., AND Y. SANNIKOV (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104(2), 379–421.

BRYZGALOVA, S., AND C. JULLIARD (2015): "The Consumption Risk of Bonds and Stocks," mimeo, LSE.

BURNSIDE, C. (2011): "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Comment," *American Economic Review*, 101(7), 3456–76.

CAMPBELL, J. Y. (1996): "Understanding Risk and Return," *Journal of Political Economy*, 104(2), 298–345.

CAMPBELL, J. Y., AND R. J. SHILLER (1988): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1(3), 195–228.

CHAVA, S., AND A. HSU (2015): "Financial Constraints, Monetary Policy Shocks, and the Cross-Section of Equity Returns," mimeo, Georgia Institute of Technology.

CHEN, N.-F., R. ROLL, AND S. A. ROSS (1986): "Economic Forces and the Stock Market," *The Journal of Business*, 59(3), 383–403.

CHIN, M., AND C. POLK (2015): "A forecast evaluation of expected equity return measures," Bank of England working papers 520, Bank of England.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): "Monetary policy shocks: What have we learned and to what end?," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, chap. 2, pp. 65–148. Elsevier.

——— (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113(1), 1–45.

CHRISTIANO, L. J., R. MOTTO, AND M. ROSTAGNO (2014): "Risk Shocks," *American Economic Review*, 104(1), 27–65.

CLOYNE, J., AND P. HURTGEN (2016): "The macroeconomic effects of monetary policy: a new measure for the United Kingdom," *American Economic Journal: Macroeconomics*, 8(4), 75–102.

COCHRANE, J. H. (1994): "Shocks," *Carnegie-Rochester Conference Series on Public Policy*, 41(1), 295–364.

——— (1996): "A Cross-Sectional Test of an Investment-Based Asset Pricing Model," *Journal of Political Economy*, 104(3), 572–621.

——— (2005): *Asset Pricing*. Princeton University Press.

——— (2011): "Presidential Address: Discount Rates," *Journal of Finance*, 66(4), 1047–1108.

——— (2017): "Macro-Finance," *Review of Finance*, 21(3), 945–985.

COCHRANE, J. H., AND L. P. HANSEN (1992): "Asset Pricing Explorations for Macroeconomics," in *NBER Macroeconomics Annual 1992, Volume 7*, NBER Chapters, pp. 115–182. National Bureau of Economic Research, Inc.

CONSTANTINIDES, G., AND D. DUFFIE (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104(2), 219–40.

CURDIA, V., AND R. REIS (2010): "Correlated Disturbances and U.S. Business Cycles," NBER Working Papers 15774, National Bureau of Economic Research, Inc.

DIMSON, E., S. NAGEL, AND G. QUIGLEY (2003): "Capturing the Value Premium in the United Kingdom," *Financial Analysts Journal*, 59(6), 35–45.

FAMA, E., AND K. FRENCH (1988): "Dividend yields and expected stock returns," *Journal of Financial Economics*, 22(1), 3 – 25.

———— (1989): "Business conditions and expected returns on stocks and bonds," *Journal of Financial Economics*, 25(1), 23–49.

———— (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33(1), 3–56.

———— (2015): "A five-factor asset pricing model," *Journal of Financial Economics*, 116(1), 1 – 22.

———— (2016): "Dissecting Anomalies with a Five-Factor Model," *Review of Financial Studies*, 29(1), 69–103.

FAMA, E. F., AND J. D. MACBETH (1973): "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81(3), 607–36.

FAMA, E. F., AND G. SCHWERT (1977): "Asset returns and inflation," *Journal of Financial Economics*, 5(2), 115–146.

FERNALD, J. (2012): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," FRBSF Working Paper 2012-19.

FERSON, W. E., AND C. R. HARVEY (1991): "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99(2), 385–415.

FRY, R., AND A. PAGAN (2011): "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 49(4), 938–60.

GALI, J. (2008): *Monetary Policy, Inflation and the Business Cycle*. Princeton University Press.

GERTLER, M., AND P. KARADI (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7(1), 44–76.

GERTLER, M., AND N. KIYOTAKI (2010): "Financial Intermediation and Credit Policy in Business Cycle Analysis," in *Handbook of Monetary Economics*, vol. 3, chap. 11, pp. 547–599. Elsevier.

GOLUB, G. H., AND C. F. V. LOAN (1996): *Matrix Computations*. The Johns Hopkins University Press, third edn.

GOYAL, A., AND I. WELCH (2008): "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21(4), 1455–1508.

GREENWALD, D. L., M. LETTAU, AND S. LUDVIGSON (2015): "Origins of Stock Market Fluctuations," CEPR Discussion Papers 10336, C.E.P.R. Discussion Papers.

GURKAYNAK, R. S., B. SACK, AND E. SWANSON (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1(1).

GURKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2007): "The U.S. Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 54(8), 2291 – 2304.

HANSEN, L. P., J. C. HEATON, AND N. LI (2008): "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy*, 116(2), 260–302.

HANSEN, L. P., AND R. J. HODRICK (1980): "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 88(5), 829–853.

HANSEN, L. P., AND R. JAGANNATHAN (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99(2), 225–62.

——— (1997): " Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52(2), 557–90.

HANSEN, L. P., AND S. F. RICHARD (1987): "The Role of Conditioning Information in Deducing Testable," *Econometrica*, 55(3), 587–613.

HARVEY, C. R., Y. LIU, AND H. ZHU (2016): "...and the Cross-Section of Expected Returns," *The Review of Financial Studies*, 29(1), 5–68.

HE, Z., B. KELLY, AND A. MANELA (2017): "Intermediary Asset Pricing: New Evidence from Many Asset Classes," *Journal of Financial Economics*, Forthcoming.

HE, Z., AND A. KRISHNAMURTHY (2014): "A Macroeconomic Framework for Quantifying Systemic Risk," Working Paper 19885, National Bureau of Economic Research.

JEGADEESH, N., AND S. TITMAN (1993): "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance*, 48(1), 65–91.

JERMANN, U. J. (1998): "Asset pricing in production economies," *Journal of Monetary Economics*, 41(2), 257–275.

KILIAN, L., AND H. LUTKEPOHL (2016): *Structural Vector Autoregressive Analysis*. Cambridge University Press.

KILIAN, L., AND D. P. MURPHY (2012): "Why Agnostic Sign Restrictions Are Not Enough: Understanding The Dynamics Of Oil Market Var Models," *Journal of the European Economic Association*, 10(5), 1166–1188.

KOIJEN, R. S., H. LUSTIG, AND S. VAN NIEUWERBURGH (2017): "The cross-section and time series of stock and bond returns," *Journal of Monetary Economics*, 88(C), 50–69.

KURMANN, A., AND C. OTROK (2013): "News Shocks and the Slope of the Term Structure of Interest Rates," *American Economic Review*, 103(6), 2612–32.

LETTAU, M., AND S. LUDVIGSON (2001a): "Consumption, Aggregate Wealth, and Expected Stock Returns," *The Journal of Finance*, 56(3), 815–849.

——— (2001b): "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying," *Journal of Political Economy*, 109(6), 1238–1287.

——— (2010): "Measuring and Modeling Variation in the Risk-Return Trade-off," in *Handbook of Financial Econometrics: Tools and Techniques*, vol. 1 of *Handbooks in Finance*, pp. 617 – 690. North-Holland, San Diego.

LETTAU, M., AND J. A. WACHTER (2011): "The term structures of equity and interest rates," *Journal of Financial Economics*, 101(1), 90 – 113.

LEWELLEN, J., S. NAGEL, AND J. SHANKEN (2010): "A skeptical appraisal of asset pricing tests," *Journal of Financial Economics*, 96(2), 175 – 194.

LIU, L. X., AND L. ZHANG (2008): "Momentum Profits, Factor Pricing, and Macroeconomic Risk," *Review of Financial Studies*, 21(6), 2417–2448.

LUCCA, D. O., AND E. MOENCH (2015): "The Pre-FOMC Announcement Drift," *Journal of Finance*, 70(1), 329–371.

LUDVIGSON, S. C., S. MA, AND S. NG (2015): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?," Working Paper 21803, National Bureau of Economic Research.

———— (2017): "Shock Restricted Structural Vector-autoregression," mimeo, National Bureau of Economic Research.

LUDVIGSON, S. C., AND S. NG (2009): "Macro Factors in Bond Risk Premia," *Review of Financial Studies*, 22(12), 5027–5067.

LUENBERGER, D. G. (1969): *Optimization by Vector Space Methods*. John Wiley and Sons.

MAIO, P., AND P. SANTA-CLARA (2012): "Multifactor models and their consistency with the ICAPM," *Journal of Financial Economics*, 106(3), 586–613.

MEHRA, R., AND E. C. PRESCOTT (1985): "The equity premium: A puzzle," *Journal of Monetary Economics*, 15(2), 145–161.

MERTENS, K., AND M. O. RAVN (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103(4), 1212–47.

NEGRO, M. D., G. EGGERTSSON, A. FERRERO, AND N. KIYOTAKI (2017): "The Great Escape? A Quantitative Evaluation of the Fed's Liquidity Facilities," *American Economic Review*, 107(3), 824–857.

OZDAGLI, A., AND M. VELIKOV (2016): "Show Me the Money: The Monetary Policy Risk Premium," mimeo, Federal Reserve Bank of Richmond.

PARKER, J. A., AND C. JULLIARD (2005): "Consumption Risk and the Cross Section of Expected Returns," *Journal of Political Economy*, 113(1), 185–222.

PASTOR, L., AND R. F. STAMBAUGH (2009): "Predictive Systems: Living with Imperfect Predictors," *Journal of Finance*, 64(4), 1583–1628.

PETKOVA, R. (2006): "Do the Fama-French Factors Proxy for Innovations in Predictive Variables?," *The Journal of Finance*, 61(2), 581–612.

PIAZZESI, M., M. SCHNEIDER, AND S. TUZEL (2007): "Housing, consumption and asset pricing," *Journal of Financial Economics*, 83(3), 531–569.

PINTER, G. (2016): "The Macroeconomic Shock with the Highest Price of Risk," Discussion Papers 1623, Centre for Macroeconomics (CFM).

RAMEY, V. A. (2016): "Macroeconomic Shocks and Their Propagation," vol. 2 of *Handbook of Macroeconomics*, pp. 71 – 162. Elsevier.

ROLL, R. (1977): "A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory," *Journal of Financial Economics*, 4(2), 129–176.

ROMER, C. D., AND D. H. ROMER (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94(4), 1055–1084.

RUBIO-RAMIREZ, J. F., D. F. WAGGONER, AND T. ZHA (2010): "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies*, 77(2), 665–696.

RUDEBUSCH, G. D. (1998): "Do Measures of Monetary Policy in a Var Make Sense?," *International Economic Review*, 39(4), 907–931.

SHANKEN, J. (1992): "On the Estimation of Beta-Pricing Models," *Review of Financial Studies*, 5(1), 1–33.

SIMS, C. A. (1980): "Macroeconomics and Reality," *Econometrica*, 48(1), 1–48.

——— (1992): "Interpreting the macroeconomic time series facts: The effects of monetary policy," *European Economic Review*, 36(5), 975–1000.

——— (1998): "Comment on Glenn Rudebusch's "Do Measures of Monetary Policy in a Var Make Sense?"," *International Economic Review*, 39(4), 933–941.

SIMS, C. A., J. H. STOCK, AND M. W. WATSON (1990): "Inference in Linear Time Series Models with Some Unit Roots," *Econometrica*, 58(1), 113–44.

SIMS, C. A., AND T. ZHA (2006): "Were There Regime Switches in U.S. Monetary Policy?," *American Economic Review*, 96(1), 54–81.

SIMS, E. R. (2016): "Differences in Quarterly Utilization-Adjusted TFP by Vintage, with an Application to News Shocks," Working Paper 22154, National Bureau of Economic Research.

SMETS, F., AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97(3), 586–606.

STOCK, J. H., AND M. W. WATSON (2012): "Disentangling the Channels of the 2007-2009 Recession," NBER Working Papers 18094, National Bureau of Economic Research, Inc.

TAYLOR, J. B. (1993): "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy*, 39(1), 195–214.

TENREYRO, S., AND G. THWAITES (2016): "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 8(4), 43–74.

UHLIG, H. (2004): "What moves GNP?," Econometric Society 2004 North American Winter Meetings 636, Econometric Society.

——— (2005): "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, 52(2), 381–419.

WEBER, M. (2015): "Nominal Rigidities and Asset Pricing," mimeo, University of Chicago.

WOLD, H. (1954): "Causality and Econometrics," *Econometrica*, 22(2), 162–177.

WOODFORD, M. (2003): *Interest and Prices.* Princeton University Press.

ZHELEZOV, O. I. (2017): "N-dimensional Rotation Matrix Generation Algorithm," *American Journal of Computational and Applied Mathematics*, 7(2), 51–57.

Appendix

A   Additional Theoretical Results

A.1   Finding the λ-shock

Proof of Proposition 1. It suffices to find an angle $\theta^\star$ and associated rotation $r_{\theta^\star}$ such that m will be a scaled multiple of any one of the rotated vectors denoted by $f_1^\star$. If $\theta^\star$ exists then $\lambda_2^\star = 0$ because $f_1 \perp f_2$ and $r_{\theta^\star}$ is an orthonormal transformation. The angle $\theta^\star = \arctan\left(\frac{\lambda_2}{\lambda_1}\right)$ satisfies $f_1^\star = r_{\theta^\star} f_1$ so that $m = \lambda_1^\star f_1^\star + \lambda_2^\star f_2^\star$ with the associated scalars $\lambda_1^\star = \frac{\|m\|}{\|f_1^\star\|} = \|m\|$ and $\lambda_2^\star = 0$. ∎

A.2   Equivalence between β-pricing Models and Linear Models of the SDF

Theorem 7 (Cochrane 2005) *Denoting the SDF, the pricing factor, the excess returns and the first- and second-stage regression coefficients from a linear pricing model by m, f, R^e, β and λ, respectively, and given the model*

$$m = 1 + [f - \mathbb{E}(f)]' b$$
$$0 = \mathbb{E}(mR^e), \tag{A.1}$$

one can find λ such that

$$\mathbb{E}(R^e) = \beta'\lambda, \tag{A.2}$$

where β are the multiple regression coefficients of excess returns R^e on the factors. Conversely, given λ in A.2, we can find b such that A.1 holds.

**Proof of Theorem 7.** See p. 106-107 of Cochrane (2005). ∎

Cochrane (2005) shows that λ and b are related $\lambda = -var(f) b$. This result simplifies greatly when working with pricing factors (such as orthonormalised VAR residuals) that have zero mean and unit variance. In this case, $\lambda = -b$ and $\mathbb{E}(f) = 0$. As a result of the linearity of the pricing model and the linearity of the VAR model, finding the orthonormalised shock in a VAR of any dimension that demands the highest price of risk (λ) when pricing a given portfolio of assets is equivalent to finding a single time series that is a linear combination of the reduced form innovations of the VAR which summarises all the information relevant to pricing the given portfolio.[22]

---

[22]Another way of saying this is that the cross-sectional R^2-measure associated with a pricing model that includes all the reduced-form residuals from the VAR is the same as the R^2-measure associated with the one-factor model which uses the appropriately orthonormalised shock. This will be confirmed during the empirical application of the method (Panel A and B of Tables 6–9.)

A.3   The Geometry of the λ-shock

To highlight the geometric nature of the orthonormalisations method and the interplay between the VAR model and the test assets, I illustrate the relevant mathematical background in a three-dimensional graph (Figure 7). There is an underlying probability space, and L_2 denotes the collection of random variables with finite variances defined on that space. L_2 is a Hilbert-space with the associated norm $\|p\| = \left(E\left(p^2\right)\right)^{1/2}$ for $p \in L_2$. Let P denote the space of portfolio excess returns (zero-price payoffs) that is a closed linear subspace of L^2.[23] P is represented by the red plane in Figure 7. An admissible SDF is a random variable m in L^2 such that $0 = E\left(mp\right)$ for all $p \in P$. The set of all admissible SDFs denoted by M is represented by the black line perpendicular to the red plane.[24]

Figure 7: A Simplified Geometry of Constructing the λ-shock
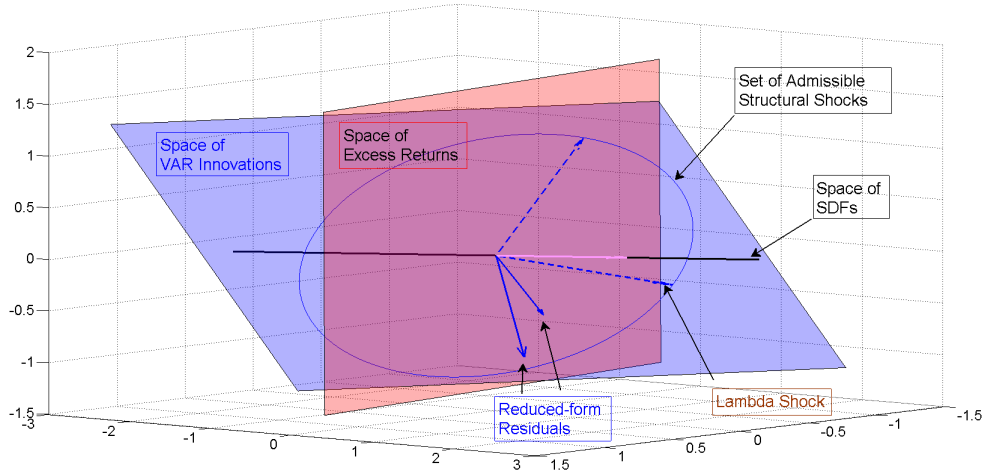
Let S denote the set of reduced-form VAR innovations (the blue solid arrows) and denote D the space spanned by these innovations. D is assumed to be a closed subspace of L_2, and it is represented by the blue plane in the Figure. The Gram-Schmidt orthogonalisation procedure allows the reduced-form innovations that span D to be transformed into a set of orthonormal vectors that also span D. The blue dashed arrows in Figure 7 represent two possible elements of the infinite sequence of orthonormalisations. The set of all admissible orthonormalisations is denoted by O and is represented by the blue circle with unit radius in the Figure.

The space of VAR innovations is unlikely to contain an SDF because of model misspecification or measurement error associated with observing SDFs (Roll, 1977). Loosely speaking, the tilted nature of the blue plane prevents all elements of O to be orthogonal

---

[23]See Hansen and Jagannathan (1991, 1997) for a detailed discussion.

[24]As is well known, all SDFs can be represented as the sum of the minimum norm SDFs (the intersection of the black line and the red plane in Figure 7) and of a random variable that is orthogonal to the space P of excess returns (Hansen and Richard (1987); Cochrane (2005)).

to the space of excess returns, i.e. $M \cap O = \emptyset$. Yet, one can find an element in O that is closest to M in the spirit of Hansen and Jagannathan (1997) by applying the classical Projection Theorem.[25] This implies that the reduced-form VAR residuals induce *one* particular orthonormal shock, which is closer to the SDF than all the other orthonormal shocks. This is the blue arrow labelled as the λ-shock in Figure 7, whose projection onto the space of SDFs is the magenta line. This shock is the best possible approximation of the SDF: it summarises all the relevant information contained in all the reduced-form residuals of the VAR, i.e. in the blue plane.

Figure 7 highlights how the modelling of the macroeconomic dynamics (the VAR innovation space from which I recover the SDF) is somewhat disjointed from the modelling of the cross-section of asset prices (the space that induces the SDF), as mentioned in the Introduction. The link between the two spaces is the orthogonality condition $0 = E(mp)$. Changing the test portfolios can be thought of as tilting the red plane while fixing the blue plane in Figure 7. In turn, augmenting the VAR with additional state variables in order to better explain/price the given test assets can be thought of as tilting the blue plane while fixing the red plane in the Figure. For example, a VAR with good (bad) pricing performance would imply a flatter (steeper) blue plane with respect to the red plane.

A.4 Additional Results

Lemma 8 *Suppose the SDF is modelled in an unconditional asset pricing framework as linear combinations of orthogonalised shocks from a VAR. The estimated prices of risk are dependent on identifying assumptions about $\boldsymbol{B}^\star$, but the estimated time-series of the SDF is independent of them.*

This statement highlights that orthogonalised shocks in a VAR are merely different linear combinations of the reduced-form residuals, thereby containing the same information set as the reduced-form innovation when pricing the cross-section of returns. In the language of empirical asset pricing: assumptions about VAR identification determine risk exposures and factor risk premia, but they do not affect the overall cross-sectional (R^2-type) fit of the transformed residuals, if all the orthogonalised shocks were to be used for pricing the cross-section of returns.

Proof of Lemma 8. The proof proceeds in three simple steps: (i) I apply arbitrary identifying assumptions to obtain a set of orthogonalised shocks (ii) I derive the estimator of the price of risk associated with the orthogonalised shocks as pricing factors (iii) and show that the implied SDF is independent of the identifying assumptions.

---

[25]That is, assuming that O is a complete linear subspace of H, there exists a unique vector $m_0 \in O$, corresponding to any vector $x \in M$, such that $\|x - m_0\| \leq \|x - m\|$ for all $m \in O$. See pp. 50-51 of Luenberger (1969) for a classic treatment and pp. 608-609 of Hansen and Richard (1987) for a conditional version of the theorem.

Let $\overline{Y}$ be an $1 \times n$ vector of average excess returns, $\tilde{Y}$ is a $T \times n$ matrix of demeaned time-series of excess returns, and η is a $T \times k$ matrix of reduced-form residuals from a k-variable VAR of any order with variance-covariance matrix Ω. Apply Cholesky decomposition to obtain triangularised innovations as pricing factors $Z = \eta \left((\boldsymbol{B})^{-1} \right)' = \eta \left((chol(\Omega))^{-1} \right)'$. The estimated risk exposures are given by the first-stage βs from time-series regressions:

$$\hat{\beta} = (Z'Z)^{-1} Z'\tilde{Y} \quad .$$

To estimated prices of risk are obtained by the second-stage cross-sectional regression:

$$
\begin{aligned}
\hat{\lambda} &= (\beta\beta')^{-1} \beta\overline{Y}' \\
&= \left[(Z'Z)^{-1} Z'\tilde{Y} \left(\tilde{Y}'Z \right) \left((Z'Z)^{-1} \right) \right]^{-1} (Z'Z)^{-1} Z'\tilde{Y}\overline{Y}' \\
&= (Z'Z) \left(Z'\tilde{Y} \left(\tilde{Y}'Z \right) \right)^{-1} Z'\tilde{Y}\overline{Y}'.
\end{aligned}
\tag{A.3}
$$

Express the reduced-form innovations in terms of orthogonalised shocks, $Z = \eta \left((\boldsymbol{B})^{-1} \right)' \equiv \eta\Delta$ re-write A.3:

$$
\begin{aligned}
\hat{\lambda} &= (\Delta'\eta'\eta\Delta) \left(\Delta'\eta'\tilde{Y} \left(\tilde{Y}'\eta\Delta \right) \right)^{-1} \Delta'\eta'\tilde{Y}\overline{Y}' \\
&= \Delta' \left(\eta'\eta \right) \Delta \left(\Delta \right)^{-1} \left(\left(\eta'\tilde{Y}\tilde{Y}'\eta \right) \right)^{-1} \left(\Delta' \right)^{-1} \Delta'\eta'\tilde{Y}\overline{Y}' \\
&= \Delta' \left(\eta'\eta \right) \left(\eta'\tilde{Y}\tilde{Y}'\eta \right)^{-1} \eta'\tilde{Y}\overline{Y}',
\end{aligned}
\tag{A.4}
$$

which proves that the estimated prices of risk depend on $\Delta \equiv \left((\boldsymbol{B})^{-1} \right)'$ which in turn depends on the identifying assumptions imposed on the structural impact matrix $\boldsymbol{B}$. The implied linear model for the SDF is written as:

$$
\begin{aligned}
m &= Z\hat{\lambda} \\
&= \eta\Delta\Delta' \left(\eta'\eta \right) \left(\eta'\tilde{Y}\tilde{Y}'\eta \right)^{-1} \eta'\tilde{Y}\tilde{Y}'\overline{Y}' \\
&= \eta\Omega \left(\eta'\eta \right) \left(\eta'\tilde{Y}\tilde{Y}'\eta \right)^{-1} \eta'\tilde{Y}\tilde{Y}'\overline{Y}',
\end{aligned}
\tag{A.5}
$$

which shows that the implied SDF depends on the *reduced-form* variance covariance matrix, Ω, and does not depend on orthogonalisation assumptions. $\blacksquare$

Proof of Proposition 3. The proof proceeds in three simple steps: (i) I derive a general form of the structural impact matrix $\boldsymbol{B}^{\star}$ and its inverse $\boldsymbol{B}^{\star-1}$ without any reference to asset pricing; (ii) I use the linear model of the SDF (2.2) to express the elements in the row of $\boldsymbol{B}^{\star-1}$ corresponding to the λ-shock, i.e. this row determines the linear relationship between the reduced form residuals and SDF innovations; finally (iii) I match the values obtained in step (i)-(ii).

Step 1 Apply the Cholesky algorithm to the reduced form variance covariance matrix $\Omega = \begin{bmatrix} \omega_{11} & \omega_{12} \\ \omega_{12} & \omega_{22} \end{bmatrix}$ to obtain a candidate for $\boldsymbol{B}$. Because Ω is positive definite, $\boldsymbol{B}$ exists and can be written as:

$$\boldsymbol{B} = chol\left(\Sigma\right) = \begin{bmatrix} \sqrt{\omega_{11}} & 0 \\ \frac{\omega_{12}}{\sqrt{\omega_{11}}} & \sqrt{\omega_{22} - \left(\frac{\omega_{12}}{\sqrt{\omega_{11}}}\right)^2} \end{bmatrix}. \tag{A.6}$$

It is known (Fry and Pagan, 2011) that one can take any orthonormal matrix Q to obtain a new structural impact matrix $\boldsymbol{B}^\star = \boldsymbol{B}Q$, with the associated set of orthogonalised shocks $e_t^\star = \eta_t \left(\boldsymbol{B}^{\star-1}\right)'$, which conforms to the reduced-form variance covariance matrix, i.e. $\Omega = \boldsymbol{B}^\star \left(\boldsymbol{B}^\star\right)' = \boldsymbol{B}Q\left(\boldsymbol{B}Q\right)' = \boldsymbol{B}Q\left(Q'\boldsymbol{B}'\right) = \boldsymbol{B}\boldsymbol{B}'$. Let Q be a rotation $r_\theta = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$, so $\boldsymbol{B}^\star = \boldsymbol{B}Q = \boldsymbol{B}r_\theta$ implies:

$$\boldsymbol{B}^\star = \begin{bmatrix} \sqrt{\omega_{11}}\cos\left(\theta\right) & -\sqrt{\omega_{11}}\sin\left(\theta\right) \\ \frac{\omega_{12}}{\sqrt{\omega_{11}}}\cos\left(\theta\right) + \xi\sin\left(\theta\right) & -\frac{c}{\sqrt{\omega_{11}}}\sin\left(\theta\right) + \cos\left(\theta\right)\xi \end{bmatrix}, \tag{A.7}$$

where $\xi = \sqrt{\omega_{22} - \omega_{12}^2/\omega_{11}}$. Matrix inversion yields:

$$\boldsymbol{B}^{\star-1} = \frac{1}{\omega_{11}\psi} \begin{pmatrix} -\omega_{12}\sin(\theta) + \sqrt{\omega_{11}}\cos(\theta)\,\psi & \omega_{11}\sin(\theta) \\ -\omega_{12}\cos(\theta) - \sqrt{\omega_{11}}\sin(\theta)\,\psi & \omega_{11}\cos(\theta) \end{pmatrix}, \tag{A.8}$$

where $\psi \equiv \omega_{11}\sqrt{\frac{\omega_{11}\omega_{22} - \omega_{12}^2}{\omega_{11}}}$.

Step 2 The linear model of the SDF (2.2) can be re-written in terms of the reduced form residuals, $\eta_t = [\eta_{1t}|\eta_{2t}]$, by using the identity $f_t = [f_{1t}|f_{2t}] = \eta_t \boldsymbol{B}^{-1} = \eta_t \left(chol\left(\Omega\right)\right)^{-1}$ and the definition A.6:

$$\begin{aligned} m_t &= \lambda_1 f_{1t} + \lambda_2 f_{2t} \\ &= \left(\lambda_1 \frac{1}{\sqrt{\omega_{11}}} - \lambda_2 \frac{\omega_{12}}{\psi}\right)\eta_{1t} + \left(\lambda_2 \frac{1}{\psi}\right)\eta_{2t}. \end{aligned} \tag{A.9}$$

Applying proposition 1 implies that the SDF can be expressed by a single orthogonalised shock, $e_{1t}^\star$, where $e_t^\star = [e_{1t}^\star|e_{2t}^\star] = \eta_t \boldsymbol{B}^{\star-1}$:

$$\begin{aligned} m_t &= \lambda_1 f_{1t} + \lambda_2 f_{2t} \\ &= \left(\sqrt{\lambda_1^2 + \lambda_2^2}\right)e_{1t}^\star + 0 \cdot e_{2t}^\star. \end{aligned} \tag{A.10}$$

[Note that designating the λ-shock to be the first column of $e_t^\star$ is arbitrary, but this does not play a role given the orthogonality of the columns of $e_t^\star$.] Hence A.9 together with A.10 determines the first row of $\boldsymbol{B}^{\star-1}$ written as:

$$\boldsymbol{B}_{1,1:1,2}^{\star-1} = \begin{bmatrix} \frac{\lambda_1 \frac{1}{\sqrt{\omega_{11}}} - \lambda_2 \frac{\omega_{12}}{\psi}}{\sqrt{\lambda_1^2 + \lambda_2^2}} & \frac{\lambda_2 \frac{1}{\psi}}{\sqrt{\lambda_1^2 + \lambda_2^2}} \end{bmatrix}. \tag{A.11}$$

Step 3   Matching values of the top right elements of A.8 and A.11 yields:

$$\theta = \arcsin\left(\frac{\lambda_2}{\sqrt{(\lambda_1^2 + \lambda_2^2)}}\right).$$

∎

A.5   Rotation Matrices

To select matrix Q in a n-variable VAR model (Subsection 2), one needs to span the n-dimensional space of rotations. See Golub and Loan (1996) for a textbook treatment and Zhelezov (2017) for a recent algorithm to generate n-dimensional rotation matrices. As an example, consider the case of a four-variable VAR model, I write Q as the product of three auxiliary Givens matrices:

$$Q = Q_1 \times Q_2 \times Q_3, \tag{A.12}$$

where:

$$Q_1 = \begin{bmatrix} \cos(\theta_1) & -\sin(\theta_1) & 0 & 0 \\ \sin(\theta_1) & \cos(\theta_1) & 0 & 0 \\ 0 & 0 & \cos(\theta_2) & -\sin(\theta_2) \\ 0 & 0 & \sin(\theta_2) & \cos(\theta_2) \end{bmatrix}$$

$$Q_2 = \begin{bmatrix} \cos(\theta_3) & & -\sin(\theta_3) & 0 \\ 0 & \cos(\theta_4) & 0 & -\sin(\theta_4) \\ \sin(\theta_3) & 0 & \cos(\theta_2) & 0 \\ 0 & \sin(\theta_4) & 0 & \cos(\theta_4) \end{bmatrix}$$

$$Q_3 = \begin{bmatrix} \cos(\theta_5) & & & -\sin(\theta_5) \\ 0 & \cos(\theta_6) & -\sin(\theta_6) & 0 \\ & \sin(\theta_6) & \cos(\theta_6) & 0 \\ \sin(\theta_5) & 0 & 0 & \cos(\theta_5) \end{bmatrix}.$$

The six Euler-angles $\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6$ are then chosen appropriately so that the objective function is satisfied. A similar construction can be used for higher dimensions. The general formula for even number of dimensions, $k \geq 4$, is written as:

$$\frac{1}{2}k \prod_{s=2}^{k/2} (2s - 1), \tag{A.13}$$

which implies that the number of angles needed to span the space rapidly increases as we add more variables to the VAR. For example, while a 4-variable VAR requires merely six angles (A.12), an 8-variable VAR requires 420 angles, whereas a 10-variable VAR requires 4725 angles to be found.

A.6 Mis-measured Factors and Pricing Performance

This section presents a Monte-Carlo exercise to illustrate why using well-known identified macroeconomic shocks to price the cross-section of returns may lead to the rejection of these shocks as valid pricing factors, even though these shocks may in fact be correlated with the true SDF innovations. As mentioned in the main text, identification of macroeconomic shocks may suffer from overly restrictive identifying assumption and from mis-measurement of macroeconomic data. As a consequence, mis-measured candidates for SDF innovations can lead to drastic deterioration in pricing performance. This is the justification behind the somewhat reverse direction taken in this paper, whereby I start with asset prices and then work "backwards".

To illustrate these measurement problems, I first take n test assets to construct the SDF, $x^\star$, from the corresponding payoff space (Chapter 4 of Cochrane (2005)).[26] As test assets, I use the FF55 for the sample period 1963Q3-2015Q3 as in my baseline analysis. I then define the distorted SDF, $\tilde{x}$, by introducing a noise term, ε_t:

$$\tilde{x} = x^\star + \varepsilon_t, \quad \varepsilon_t \sim N\left(0, \sigma^2\right) \tag{A.14}$$

where σ is the standard deviation of the measurement error ε_t. To assess the pricing performance of the distorted SDF, I first estimate n time series regressions, $R_{it} = \tilde{x}\beta_i + \epsilon_{it}$, $i = 1 \ldots n$. Second , I estimate a cross-section regression, $\bar{R}_i = \tilde{\beta}_i \times \lambda + \alpha_i$, where $\bar{R}_i = \frac{1}{T}\sum_{t=1}^{T} R_{it}$, $\tilde{\beta}_i$ is the OLS estimate obtained in the first stage and α_i is a pricing error. The model's fit is then assessed using the following statistic (Burnside, 2011):

$$R^2 = 1 - \frac{\left(\bar{R} - \tilde{\beta}\tilde{\lambda}\right)'\left(\bar{R} - \tilde{\beta}\tilde{\lambda}\right)}{\left(\bar{R} - \widehat{R}\right)'\left(\bar{R} - \widehat{R}\right)}, \tag{A.15}$$

where $\widehat{R} = \frac{1}{n}\sum_{i=1}^{n} \bar{R}_i$ is the cross-sectional average of the mean returns in the data. Moreover, I compute the correlation between the true and distorted SDFs:

$$\rho^\star = corr\left(x^\star, \tilde{x}\right). \tag{A.16}$$

Then I explore how the cross-sectional fit A.15 and the correlation coefficient A.16 change as I add more noise to the SDF, i.e. I aim to estimate the derivatives $\partial R^2/\partial \sigma$ and $\partial \rho^\star/\partial \sigma$.

For the Monte-Carlo exercise, I use a grid $\sigma = [0 : 0.05 : 0.75]$ to control for the amount measurement error, and for each value of σ, I generate 5000 time-series of $\tilde{x}$, and compute the statistics using A.15 and A.16. Figure 8 shows the median values (solid lines) of the statistics together with 10-90% simulation bands (shaded areas).

---

[26]Specifically, I follow Section 4.1. of Cochrane (2005) and construct the discount factor $x^\star$ from the payoff space using $x^\star = p'E\left(xx'\right)^{-1}x$, where x denotes the test assets with payoffs p.

Figure 8: Role of Mis-measurement in Macroeconomic Shocks: Results from a Monte-Carlo Exercise

Pricing Performance of Noisy SDFs, Implied by FF55
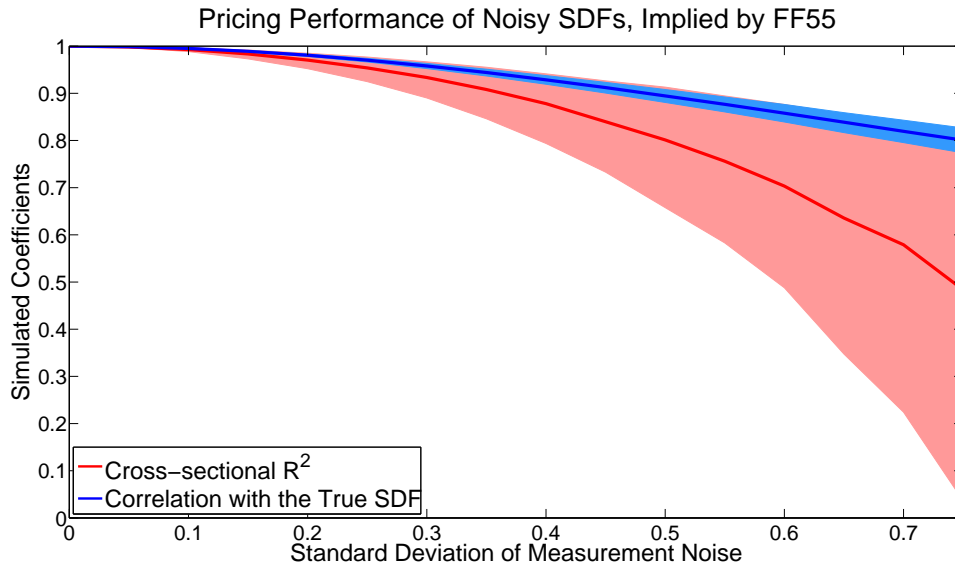
Notes: The Figure illustrates how adding noise to the SDF (constructed from the space of test assets, following 4.1. of Cochrane (2005)) changes the correlation with the true SDF (blue line) and the cross-sectional fit (red line). The cross-sectional fit is measured using the R^2 measure A.15, and the correlation measure is based on A.16. The shaded areas correspond to the 10-90% bands based on 5000 Monte-Carlo simulations of A.14. The construction of the SDF is based on the FF55 portfolios, covering the sample period is 1963Q3-2015Q3.

The results show that adding noise to the SDF deteriorates the pricing performance more quickly than it reduces the correlation between the noisy and true SDFs. Importantly, the uncertainty around the estimated R^2 increases much more rapidly than the uncertainty around the estimated $\rho^\star$. For example, for $\sigma = 0.75$, the correlation between distorted SDF and the true SDF is still above 80%, whereas the cross-sectional fit of the corresponding pricing model can result in close to zero explanatory power.

These results provide a justification for (i) why using noisy estimates of macroeconomic shocks (identified by the macroeconomics literature) directly in asset pricing tests may lead the rejection of these shocks as valid pricing factors, and (ii) why the reverse direction taken in this paper may be more successful in uncovering the empirical linkages between business cycle fluctuations and asset prices.

B Additional Empirical Results

B.1 *Robustness to Alternative Lag Structure of the VAR*

Figure 9: Impulse Responses to a λ-shock: Robustness to Lag Structure

(a) Sample: 1963Q3-2008Q3
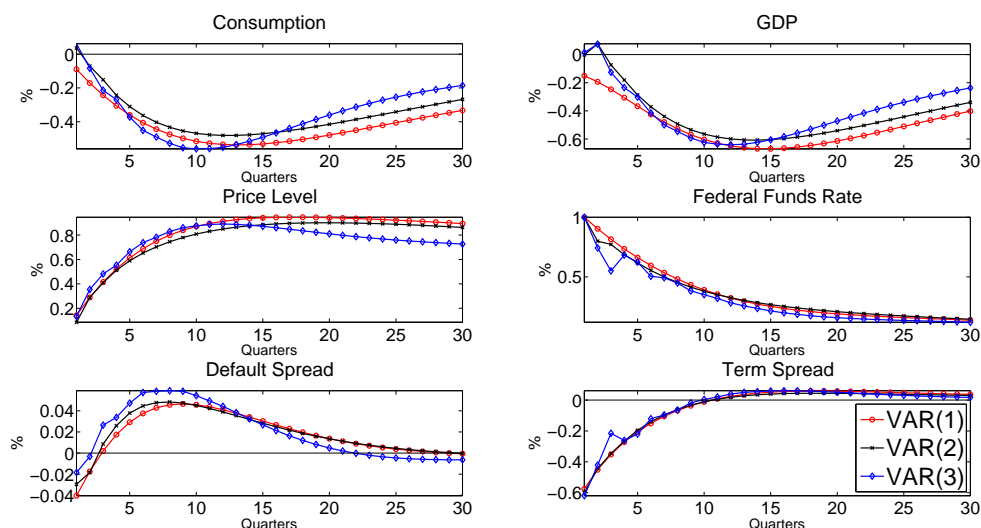
(b) Sample: 1963Q3-2015Q3
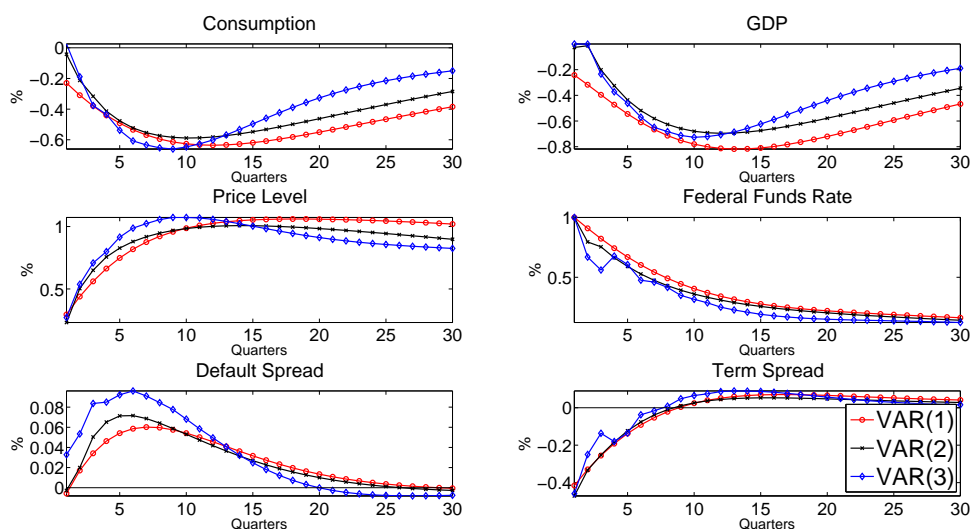
Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The IRFs are computed from VAR(1), VAR(2) and VAR(3) models. In all cases, the FF55 portfolios were used as test assets, and the IRFs are normalised to increase the federal funds rate by 100bp.

Table 3: Forecasting Excess Returns: Results from a <u>VAR(3)</u> Model

| | \multicolumn{8}{c}{Forecast Horizon H} | | | | | | | |
	1Q	2Q	3Q	4Q	5Q	6Q	7Q	8Q
\multicolumn{9}{c}{Model A: CAY}								
CAY	0.55	1.12	1.77	2.52	3.17	4.00	4.92	5.77
	(2.18)	(2.25)	(2.33)	(2.46)	(2.59)	(3.00)	(3.58)	(4.25)
	[0.01]	**[0.03]**	**[0.05]**	**[0.09]**	**[0.11]**	**[0.14]**	**[0.19]**	**[0.23]**
\multicolumn{9}{c}{Model B: Counterfactual VAR Variables Induced by the γ-Shock}								
FFR	0.47	1.00	1.50	1.87	2.24	2.75	3.31	3.80
	(1.14)	(1.19)	(1.37)	(1.45)	(1.56)	(1.77)	(2.03)	(2.29)
DEF	-7.76	-14.53	-20.40	-24.42	-28.36	-32.55	-37.35	-41.47
	(-1.84)	(-1.68)	(-1.91)	(-2.04)	(-2.19)	(-2.42)	(-2.70)	(-2.94)
TERM	3.01	6.41	9.72	12.40	14.99	17.87	20.81	23.16
	(2.71)	(2.72)	(2.90)	(2.96)	(3.09)	(3.27)	(3.53)	(3.87)
	[0.04]	**[0.10]**	**[0.17]**	**[0.23]**	**[0.28]**	**[0.33]**	**[0.38]**	**[0.42]**
\multicolumn{9}{c}{Model C: Counterfactual VAR Variables Induced by All Other Shocks}								
FFR	-0.18	-0.28	-0.25	-0.21	-0.09	0.11	0.32	0.50
	(-0.70)	(-0.64)	(-0.43)	(-0.30)	(-0.11)	(0.11)	(0.25)	(0.34)
DEF	1.82	3.71	3.91	3.71	2.94	2.27	1.32	0.27
	(0.90)	(1.23)	(0.97)	(0.79)	(0.57)	(0.39)	(0.18)	(0.03)
TERM	-0.22	-0.64	-0.58	-0.45	-0.28	0.17	0.38	0.96
	(-0.37)	(-0.58)	(-0.36)	(-0.23)	(-0.12)	(0.07)	(0.14)	(0.32)
	[-0.01]	**[-0.00]**	**[-0.00]**	**[-0.01]**	**[-0.01]**	**[-0.01]**	**[-0.01]**	**[-0.01]**

Notes: The table reports results from regressions of excess returns on lagged variables. H denotes the return horizon in quarters. The dependent variable is the sum of H log excess returns on the CRSP based S&P Composite Index. The regressors are one-period lagged values of actual time-series of the federal funds rate (FFR), the term-spread (TERM) and the default spread (DEF) in Model A, the CAY measure of Lettau and Ludvigson (2001a,b) in Model B, and the counterfactual time-series of FFR, TERM and DEF (induced by the γ-shock) from six-variable VAR(3) and VAR(4) models, estimated over 1963Q3-2015Q3. The γ-shock is constructed so that the corresponding forecast power at the four-quarter horizon is maximised. For each of the three regressions, the table reports the OLS estimates of the regressors, the t-statistics using the Hansen and Hodrick (1980) correction (as implemented in Cochrane (2011)) are in parentheses, and adjusted R^2 statistics are in the bolded square brackets. Both the CAY measure and the counterfactual predictors are treated as known variables.

Table 4: Forecasting Excess Returns: Results from a <u>VAR(4)</u> Model

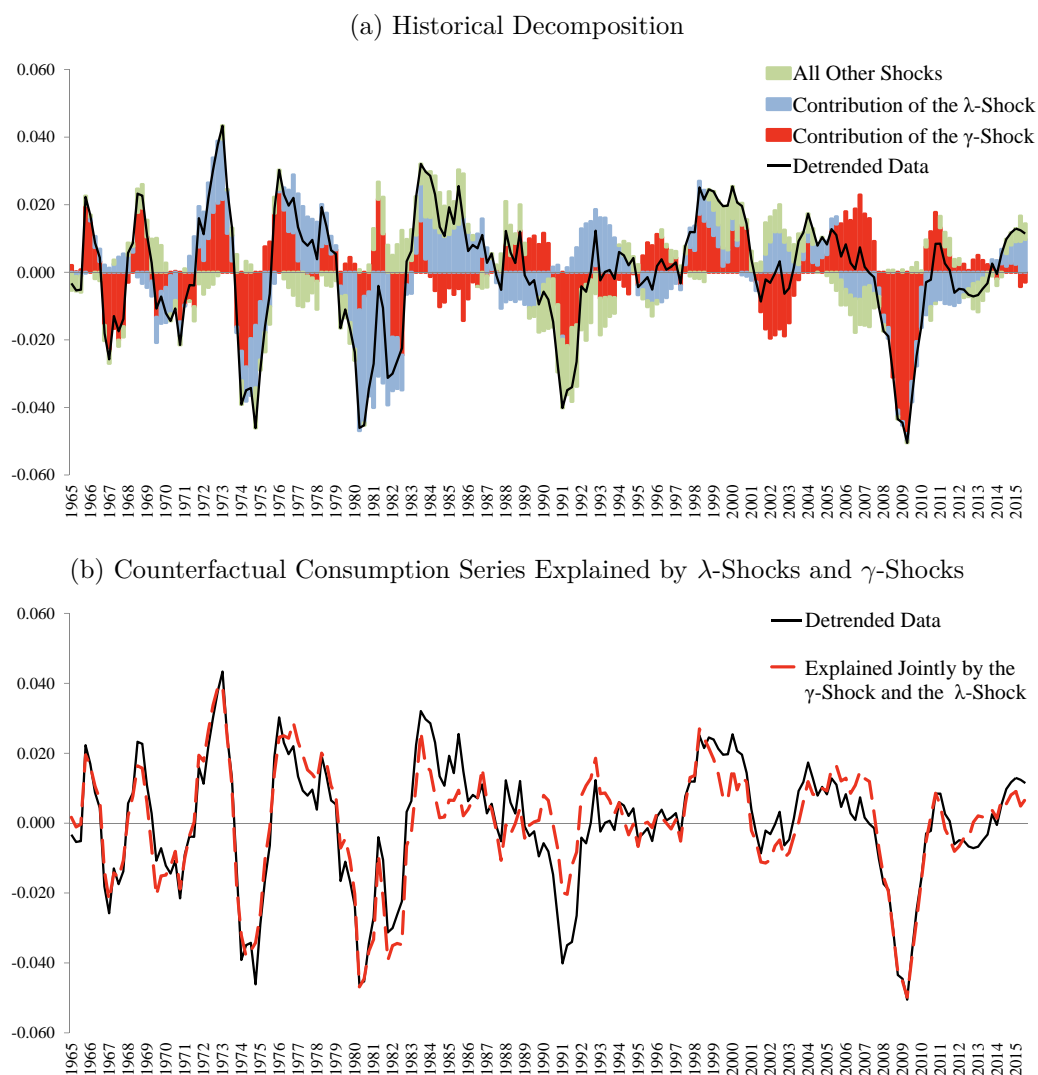|	Forecast Horizon H							
	1Q	2Q	3Q	4Q	5Q	6Q	7Q	8Q
	Model A: CAY							
CAY	0.55	1.12	1.77	2.52	3.17	4.00	4.92	5.77
	(2.18)	(2.25)	(2.33)	(2.46)	(2.59)	(3.00)	(3.58)	(4.25)
	[0.01]	**[0.03]**	**[0.05]**	**[0.09]**	**[0.11]**	**[0.14]**	**[0.19]**	**[0.23]**
	Model B: Counterfactual VAR Variables Induced by the γ-Shock							
FFR	0.40	0.74	1.07	1.35	1.64	2.02	2.47	2.98
	(1.06)	(0.94)	(1.02)	(1.09)	(1.20)	(1.39)	(1.62)	(1.92)
DEF	-7.01	-12.02	-16.29	-19.46	-22.33	-24.99	-28.18	-32.25
	(-1.88)	(-1.52)	(-1.65)	(-1.77)	(-1.94)	(-2.15)	(-2.38)	(-2.68)
TERM	2.80	5.66	8.55	11.08	13.47	16.03	18.63	21.04
	(2.71)	(2.54)	(2.68)	(2.75)	(2.90)	(3.10)	(3.38)	(3.77)
	[0.04]	**[0.09]**	**[0.15]**	**[0.20]**	**[0.25]**	**[0.30]**	**[0.34]**	**[0.39]**
	Model C: Counterfactual VAR Variables Induced by All Other Shocks							
FFR	-0.20	-0.29	-0.23	-0.18	-0.04	0.20	0.44	0.63
	(-0.78)	(-0.67)	(-0.40)	(-0.25)	(-0.04)	(0.19)	(0.34)	(0.41)
DEF	2.38	4.59	4.86	4.69	3.86	2.98	1.84	0.84
	(1.15)	(1.55)	(1.25)	(1.01)	(0.73)	(0.50)	(0.23)	(0.09)
TERM	-0.18	-0.45	-0.26	-0.05	0.20	0.79	1.13	1.77
	(-0.31)	(-0.42)	(-0.16)	(-0.03)	(0.10)	(0.37)	(0.46)	(0.65)
	[-0.00]	**[0.00]**	**[0.00]**	**[-0.00]**	**[-0.01]**	**[-0.01]**	**[-0.01]**	**[-0.00]**

Notes: The table reports results from regressions of excess returns on lagged variables. H denotes the return horizon in quarters. The dependent variable is the sum of H log excess returns on the CRSP based S&P Composite Index. The regressors are one-period lagged values of actual time-series of the federal funds rate (FFR), the term-spread (TERM) and the default spread (DEF) in Model A, the CAY measure of Lettau and Ludvigson (2001a,b) in Model B, and the counterfactual time-series of FFR, TERM and DEF (induced by the γ-shock) from six-variable VAR(3) and VAR(4) models, estimated over 1963Q3-2015Q3. The γ-shock is constructed so that the corresponding forecast power at the four-quarter horizon is maximised. For each of the three regressions, the table reports the OLS estimates of the regressors, the t-statistics using the Hansen and Hodrick (1980) correction (as implemented in Cochrane (2011)) are in parentheses, and adjusted R^2 statistics are in the bolded square brackets. Both the CAY measure and the counterfactual predictors are treated as known variables.

Table 5: The Contribution of the λ-shock and γ-shock to Business Cycles: FEV Decomposition from VAR(3) and VAR(4) Models

	VAR(3) Model							
	Consumption				Federal Funds Rate			
	λ-Shock	γ-Shock	$\lambda\&\gamma$	Unexpl.	λ-Shock	γ-Shock	$\lambda\&\gamma$	Unexpl.
1Q	0.5	72.3	72.8	27.2	64.1	3.2	67.3	32.7
2Q	3.9	72.1	76.0	24.0	54.1	9.0	63.1	36.9
3Q	10.7	70.4	81.1	18.9	51.4	11.0	62.4	37.6
4Q	15.1	66.5	81.6	18.4	50.7	14.1	64.8	35.2
8Q	32.4	44.5	76.9	23.1	47.1	20.8	67.9	32.1
16Q	43.0	25.9	68.9	31.1	43.9	24.1	68.0	32.0
32Q	42.2	20.7	62.8	37.2	42.8	24.5	67.3	32.7
	VAR(4) Model							
	Consumption				Federal Funds Rate			
	λ-Shock	γ-Shock	$\lambda\&\gamma$	Unexpl.	λ-Shock	γ-Shock	$\lambda\&\gamma$	Unexpl.
1Q	0.2	75.8	76.0	24.0	73.3	2.4	75.7	24.3
2Q	4.0	75.8	79.8	20.2	62.3	8.2	70.5	29.5
3Q	10.9	74.0	84.8	15.2	60.0	10.1	70.1	29.9
4Q	13.7	70.3	84.0	16.0	60.7	12.5	73.3	26.7
8Q	33.5	46.3	79.8	20.2	58.3	18.1	76.4	23.6
16Q	48.7	27.2	75.9	24.1	56.8	18.7	75.5	24.5
32Q	47.6	24.0	71.7	28.3	56.3	18.3	74.6	25.4

Notes: The table shows the % fraction of the total forecast error variance that is explained by the λ-shock and the γ-shock over different forecast horizons. The FF55 portfolios are used as test portfolios for the VAR models. The estimation period is 1963Q3-2015Q3.

Figure 10: Results from a $\underline{\text{VAR(4)}}$ – Decomposing Annual US Consumption Growth: the Role of λ-Shocks and γ-Shocks

(a) Historical Decomposition

(b) Counterfactual Consumption Series Explained by λ-Shocks and γ-Shocks

Notes: The figure shows the results implied by the historical decomposition from a VAR(4) estimated on the sample 1963Q3-2015Q3. I use the FF55 to construct the λ-shock, and the CRSP based SP500 return (in excess of the corresponding T-bill rate) to construct the γ-shock. The deterministic trend component, implied by the VAR, is removed from the time-series.

B.2 The λ-shock in the Monthly VAR of Uhlig (2005)

As robustness check, I investigate the price response of the λ-shock further, and use the monthly VAR of Uhlig (2005). Following his paper, I impose sign restrictions on the impulse responses of prices, nonborrowed reserves and the Federal Funds Rate in response to a monetary policy shock, thereby fixing the price puzzle anomaly while remaining agnostic about the effect of monetary policy shocks on other macrovariables of interest. The black lines and the associated error bands in Figure 11 replicate Figure 6 of Uhlig (2005), using his dataset. The blue line shows the response to a Cholesky-orthogonalised innovation in the federal funds rate that is ordered before nonborrowed and total reserves (Figure 5 of Uhlig (2005)). The red and purple lines show the responses to a λ-shock that is constructed using the FF55 and FF25 portfolios, respectively.

Figure 11: The VAR Model of Uhlig (2005): Monetary Policy Shocks and λ-shocks
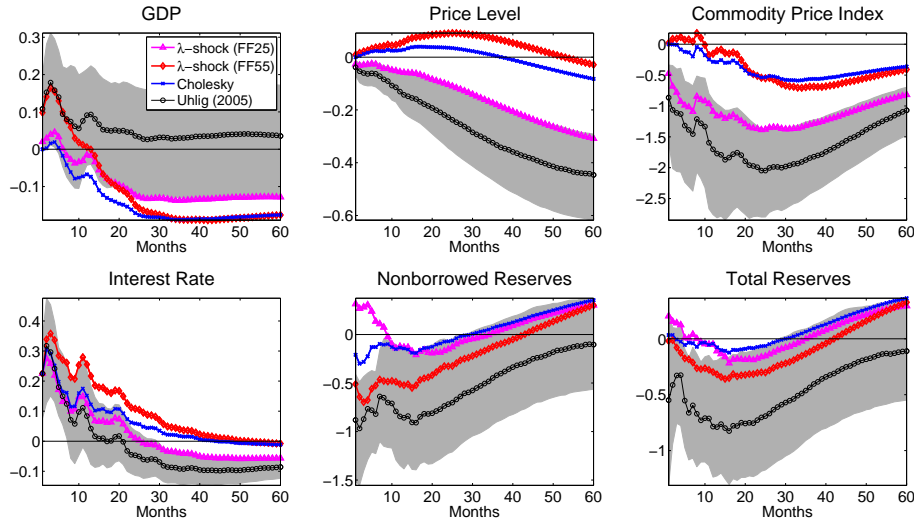
Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in months. The monthly VAR(12) is estimated over 1965m1-2003m12, using the dataset of Uhlig (2005). The black lines and associated 16-84 bands are identical to Figure 6 of Uhlig (2005) that shows the responses to a monetary policy shock identified by the 'pure sign restriction approach'. To construct the λ-shock, I use the FF25 (purple lines) and the FF55 (red lines) portfolios. The blue lines are the are Cholesky-orthogonalised (Sims, 1980) interest rate shock. The IRFs are normalised to have the same contemporaneous effect on the interest rate.

While using a different dataset at monthly (instead of quarterly) frequency, the estimated λ-shock continues to induce business cycles similar to that caused by monetary policy shocks. Note, however that the λ-shocks generate a more contractionary GDP response than the monetary policy shock identified by Uhlig (2005). However, recent papers (Antolin-Diaz and Rubio-Ramirez, 2018; Arias, Rubio-Ramirez, and Waggoner, 2018) have argued that the original identification of Uhlig (2005) retains many structural parameters with improbable implications for the systematic response of monetary policy to output, and the updated empirical evidence delivers more contractionary impulse responses of output. Overall, Figure 11 suggests that the dynamics generated by the λ-shock are in-between those generated by sign restrictions and Cholesky-orthogonalisation. The 'price puzzle' is present in the case of

the Cholesky-shock and λ-shock induced by the FF55 portfolios, but it is absent in the case of sign restrictions (by construction) and the λ-shock induced by the FF25 portfolios.

B.3 Decomposing Lower Frequency Variation in Consumption

This subsection explores the contribution of the λ-shock and the γ-shock to variation in aggregate consumption at a lower frequency. Instead of decomposing annual consumption growth (as in Figure 6 of the main text), I now decompose the deviation of the level of aggregate consumption from the trend implied by the baseline six-variable VAR model. The results are shown in Figure 12. Consistent with FEV decomposition, the λ-shock contributes to the lower frequency dynamics much more than to the higher frequency dynamics proxied by annual consumption growth.

Figure 12: Decomposing Level Deviations of US Consumption: the Role of λ-Shocks and γ-Shocks
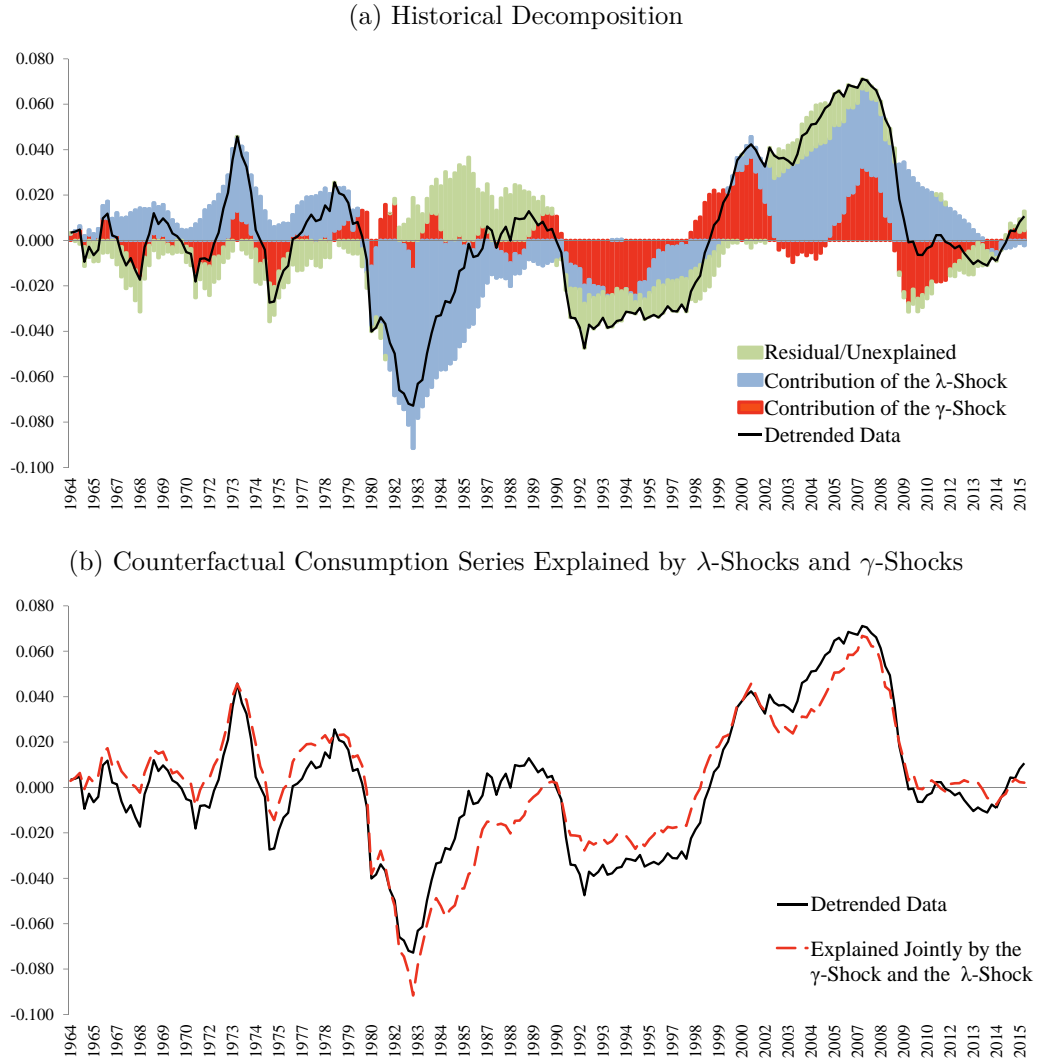
(a) Historical Decomposition

(b) Counterfactual Consumption Series Explained by λ-Shocks and γ-Shocks
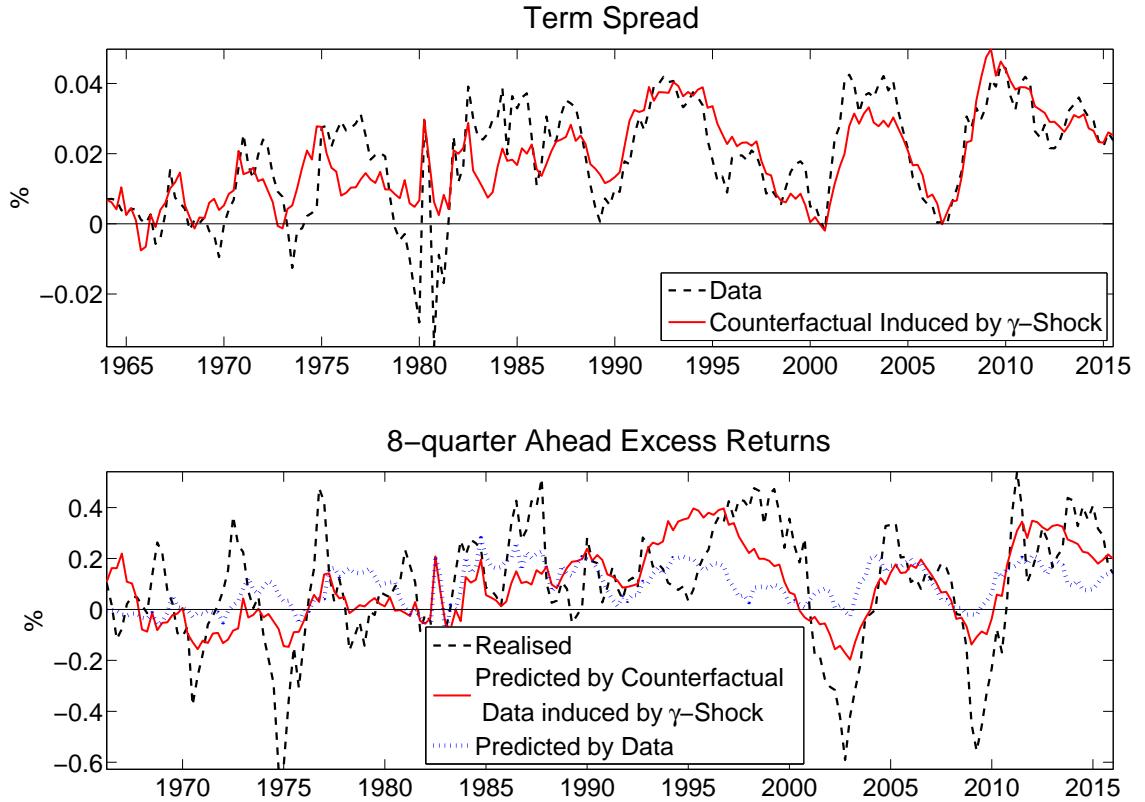

Notes: The figure shows the results implied by the historical decomposition from a VAR(2) estimated on the sample 1963Q3-2015Q3. I use the FF55 to construct the λ-shock, and the CRSP based SP500 return (in excess of the corresponding T-bill rate) to construct the γ-shock. The deterministic trend component, implied by the VAR, is removed from the time-series.

Specifically, in addition to largely affecting the early 1980s consumption decline, the λ-shock made a sizeable contribution to the persistent expansion of consumption in the 2000s above the long-run trend.

B.4 *The γ-Shock and Time-varying Risk Premia*

Figure 13: The γ-Shock Induced Counterfactual Term Spread and Forecasting Excess Returns

Notes: The upper panel of the figure shows the term spread (measured as the difference between the 10-year yield and the short-term Treasury bill rate, taken from Goyal and Welch (2008)) along with the counterfactual term spread implied by the γ-shock from a VAR(2) estimated on the sample 1963Q3-2015Q3. The lower panel shows realised 8-quarter cumulative excess returns along with the fitted values from the regression (based on 2.5) $r_{t+1}^H = a + \beta_1 FFR_t + \beta_2 DEF_t + \beta_3 TERM_t + \varepsilon_{t+1}$ with $H = 8$ (blue dotted line), and also the fitted values from the regression $r_{t+1}^H = a + \beta_1 \widehat{FFR}_t + \beta_2 \widehat{DEF}_t + \beta_3 \widehat{TERM}_t + \varepsilon_{t+1}$ with $H = 8$ (red solid line), where $\widehat{}$ denotes the counterfactual time-series implied by the γ-shock. The correlation between realised returns and the data-based predicted series and the γ-shock-based predicted series are 0.36 and 0.63, respectively.
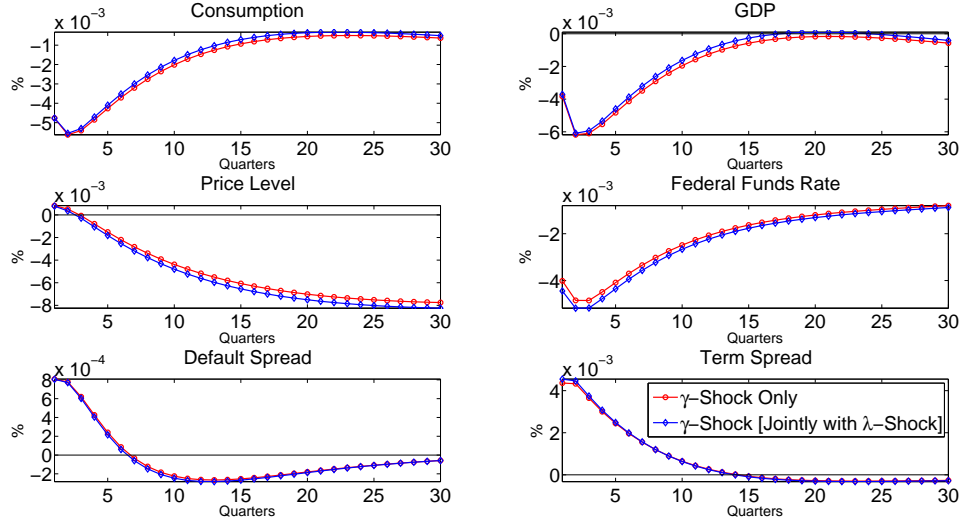
B.5 *The Orthogonality of the λ-Shock to the γ-Shock*

Figure 14: Illustrating the Orthogonality of the λ-Shock with respect to the γ-Shock

(a) The λ-Shock
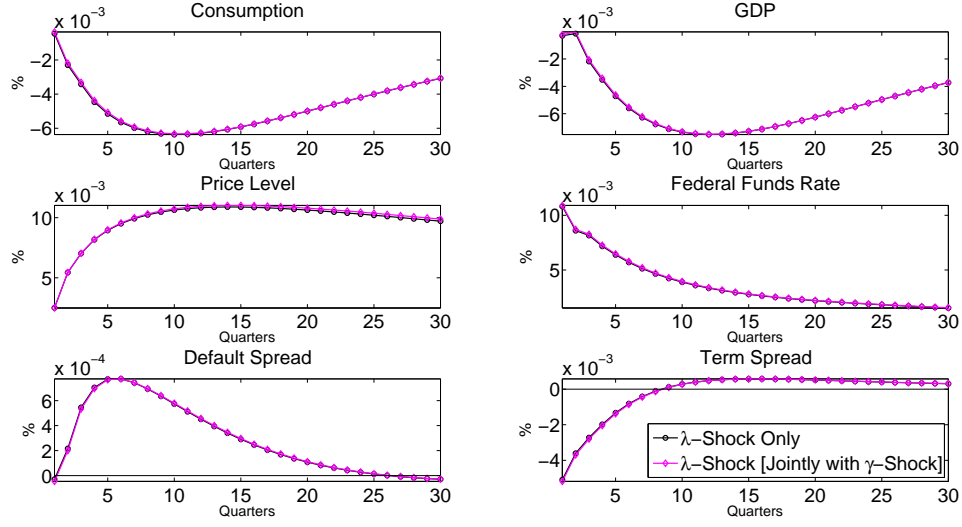
(b) The γ-Shock

Notes: The Figure illustrates the orthogonality of the λ-Shock with respect to the γ-Shock. In the upper panel, the λ-Shock is constructed with and without simultaneously constructing the γ-Shock. In the lower panel, the γ-Shock is constructed with and without simultaneously constructing the λ-Shock. The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The FF55 portfolios were used as test assets, and the sample period is 1963Q3-2015Q3.

B.6   Recession-Shocks, γ-Shock and λ−Shocks

Figure 15: Impulse Responses to a Recession-Shock, γ-Shock and to a λ-Shock
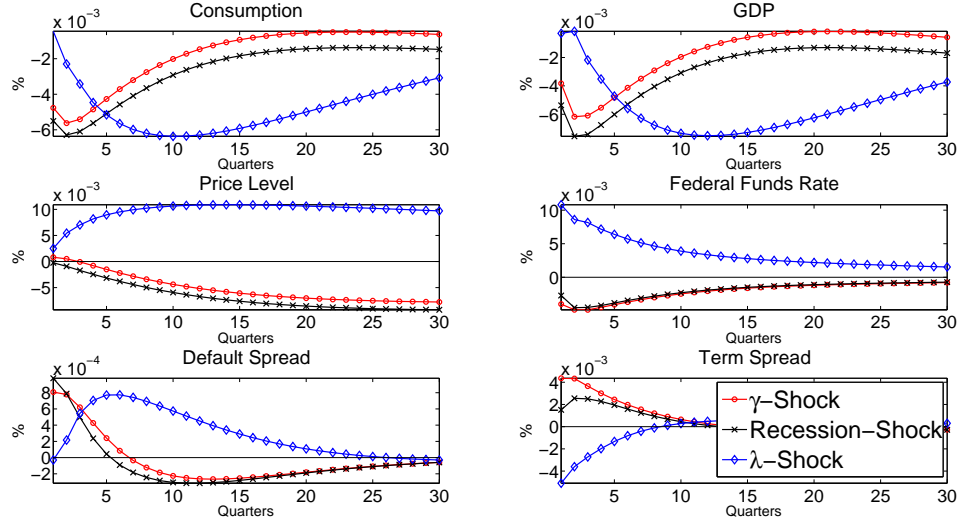
Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The VAR(2) is estimated over the 1963Q3-2015Q3 period. The shocks are 1 sd.The FF55 portfolios were used as test assets.

Figure 16: Decomposing Level Deviations of US Consumption: the Role of λ-Shocks and Recession-Shocks
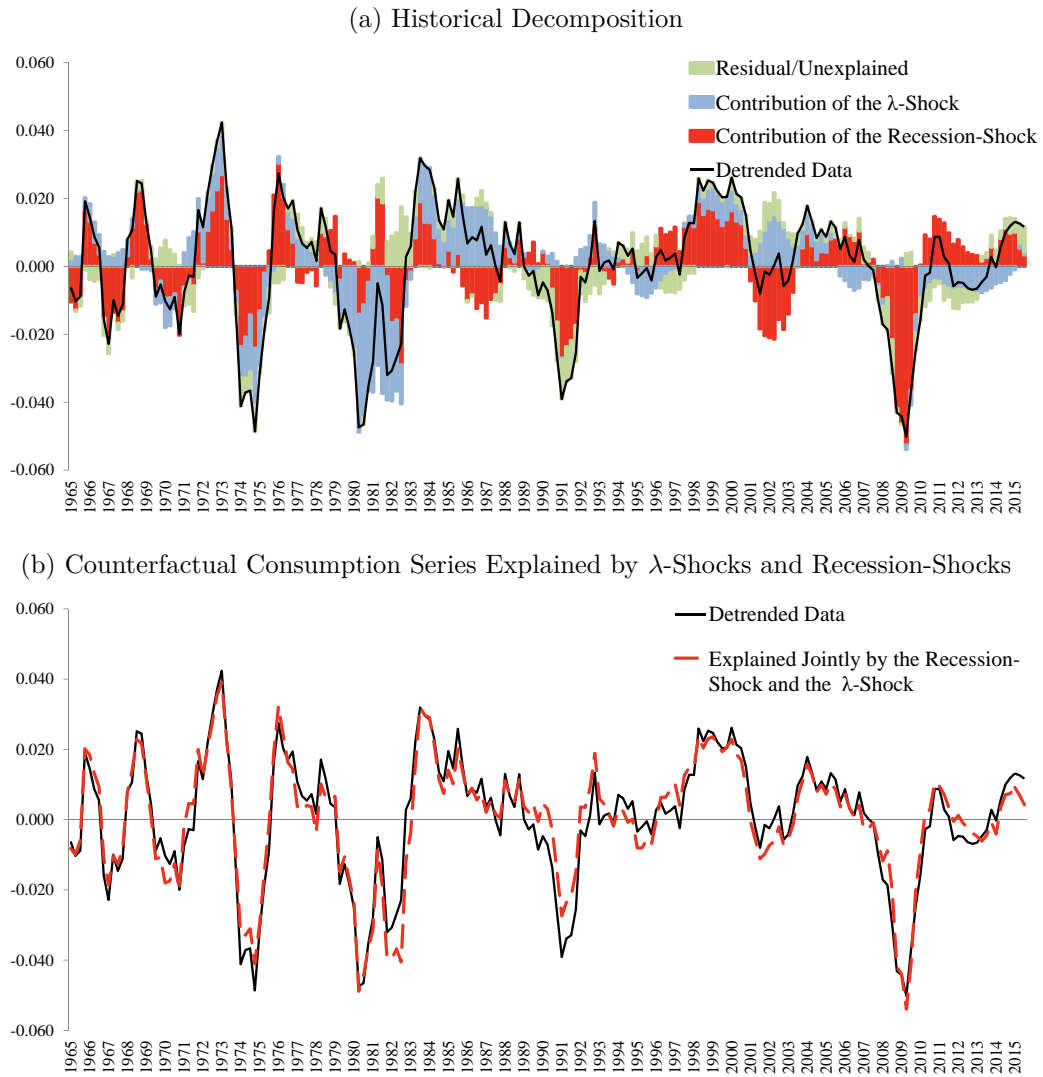
(a) Historical Decomposition

(b) Counterfactual Consumption Series Explained by λ-Shocks and Recession-Shocks

Notes: The figure shows the results implied by the historical decomposition from a VAR(2) estimated on the sample 1963Q3-2015Q3. I use the FF55 to construct the λ-shock. The maximisation of the price of risk and the minimisation of 2.7 are done jointly. The blue (red) line is the contribution of the λ-shock (Recession-shock) to the data. The purple line in the bottom panel is the sum of the contributions of the λ-shock and the Recession-shock to the data. The deterministic trend component, implied by the VAR, is removed from the time-series.

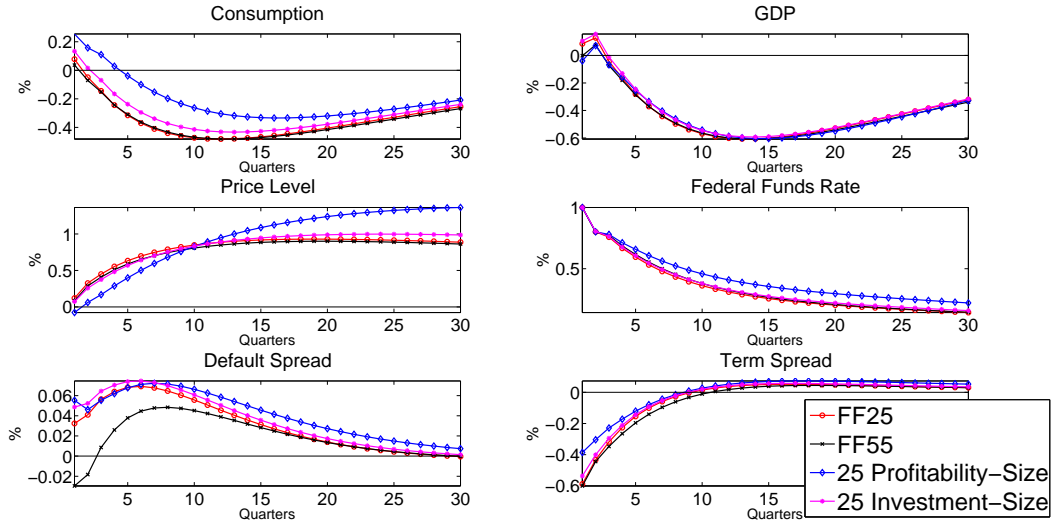## B.7   Other Equity Portfolios and Government Bond Returns

To check the robustness of the baseline results I explore how the behaviour of the λ-shock changes when the same VAR model and the orthogonalisation method are applied to other test assets. A natural choice is the 25 portfolios double sorted on size-profitability and size-investment. These portfolios feature prominently in the most recent empirical asset pricing studies (Fama and French, 2015, 2016). In addition, I also compute the IRFs for the λ-shock implied by the benchmark FF25 portfolios, sorted on size-B/M, that have been the most studied test assets to date.

The upper panel of Figure 17 shows the IRFs for these three sets of equity portfolios along the benchmark FF55 used in the main text. The results suggest that the economic behaviour of the λ-shock implied by these portfolios is very similar to each other. The only quantitative difference is that the baseline results imply a larger peak effect on consumption and a more delayed effect on the default spread compared to Figure 17.

Moreover, I also use government bond returns that are calculated using the zero coupon yield data constructed by Gurkaynak, Sack, and Wright (2007) that fit Nelson-Siegel-Svensson curves on daily data. The parameters for backing out the cross-section of yields are published on their website. The sample period is 1975Q2-2008Q3 so that I have sufficiently large cross-section of yields. I use maturities for $n = 18, 24, \ldots, 120$ months and compute one-month holding period excess returns which I then transform into quarterly series. The resulting 18 bond portfolios are used to construct the λ-shock. The lower panel of Figure 17 shows the results, confirming that the shock responsible for pricing equities is virtually identical to the shock that prices government bonds. This is consistent with the relatively small but growing literature on the joint pricing of stocks and bonds (Lettau and Wachter 2011; Bryzgalova and Julliard 2015; Koijen, Lustig, and Van Nieuwerburgh 2017).

Figure 17: Impulse Responses to a λ-shock, Implied by other Equities vs Bonds

(a) *FF25, FF55, 25 Profitability-Size and 25 Investment-Size* Portfolios

(b) *FF25* and *US Government Bond Returns*
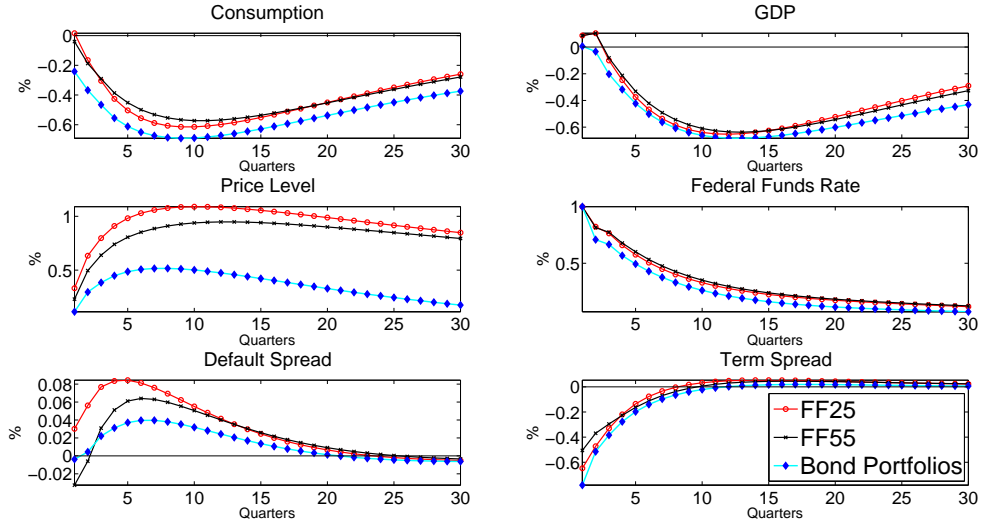
Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The upper panel uses the VAR(2), as estimated in Subsection 2, and employs alternative equity portfolios to the construction of the λ-shock. The lower panel estimates the same VAR(2) on a subsample 1975Q2-2008Q3, and constructs the λ-shock implied by the FF25 and FF25 portfolios as well as by quarterly holding period excess returns on 18 US treasury bonds with maturities $n = 18, 24, \ldots, 120$ months. I all IRFs are normalised to increase the federal funds rate by 100bp.

B.8 Adding More Variables to the VAR

One can easily add more variables to the VAR to improve the return predicting power of the γ-shock or to improve the cross-sectional pricing performance of the λ-shock.[27] This is particularly useful, given that most pricing factors (316 of them listed in Harvey, Liu, and

---

[27]Increasing the size of the VAR introduces only computational challenges. For example, in an 8-variable (10-variable) VAR one needs to find 420 (4725) angles to span the 8-dimensional space of rotations. See section A.5 of the Appendix.

Zhu (2016)) are reduced-form objects and often exhibit high correlations with one another. For example, consumption innovations and output innovations extracted from individual AR(1) models have around 66% correlation, term spread innovation and federal funds rate innovations have around -82% correlation, the intermediary capital risk factor constructed by He, Kelly, and Manela (2017) and excess returns on the market (their second pricing factor) have a 78% correlation.[28]

To study this further, I will add to the VAR the aggregate capital ratio of the financial intermediary sector (constructed by He, Kelly, and Manela (2017)). I then re-estimate the VAR on a shorter sample, 1970Q1-2015Q3 (dictated by the availability of this time-series), and calculate the dynamic effects of the λ-shock and the γ-shock on the intermediary capital ratio. Note that this procedure could be applied to any other reduced-form pricing variable of interest. Figure 18 summarises the results. Panel a shows the IRFs for a one standard deviation contractionary innovation in both shocks. In response to both shocks, the intermediary capital ratio drops immediately by about 0.3% and then gradually returns to steady-state after about five years. Note that the rest of the variables in the VAR exhibit very similar dynamics to the baseline (Figure 4). While GDP is replaced with the intermediary capital ratio in the model and this VAR is estimated on a different sample, the time-series of the λ-shock and the γ-shock have a high (around 80%) correlation across the two VAR models.

Panel b of Figure 18 shows the FEV decomposition of the intermediary capital ratio along with that of aggregate consumption. The results suggest that both orthogonalised shocks are important in driving fluctuations in the capital ratio. For example, at one-year horizon, the λ-shock and the γ-shock explain about 35% and 31% of the forecast error variance in the capital ratio, respectively. The decomposition of aggregate consumption continues to be similar to my baseline VAR (Table 2). It is to note that a considerable fraction of capital ratio fluctuations is left unexplained by the λ-shock and the γ-shock. Future work could enrich this simple VAR to increase explanatory power.

This exercise highlights that unexpected changes in the balance sheet health of financial intermediaries cannot be interpreted as purely exogenous events. While the driving force in macroeconomic models with financial intermediaries (Gertler and Kiyotaki 2010; He and Krishnamurthy 2014) is often related to exogenous movements in the capital stock ("capital quality shock"), my results show that a large fraction of the unforecastable component in the capital ratio can be explained by at least two orthogonalised macroeconomic shocks that have very different effects on business cycle fluctuations. This highlights that by purely focusing on the reduced-form unforecastable component in the capital ratio, one cannot accurately detect the nature of the macroeconomic force responsible for the observed fluc-

---

[28]These number are based on estimates of individual AR(1) models on consumption, GDP, the term spread and the Federal Funds rate covering the period 1963Q3-2008Q3. The correlation between the intermediary capital risk factor and market excess returns are for 1970Q1-2012Q4 as in He, Kelly, and Manela (2017).

Figure 18: Financial Intermediary Capital Dynamics

(a) Impulse Responses to a λ-Shock and to a γ-Shock
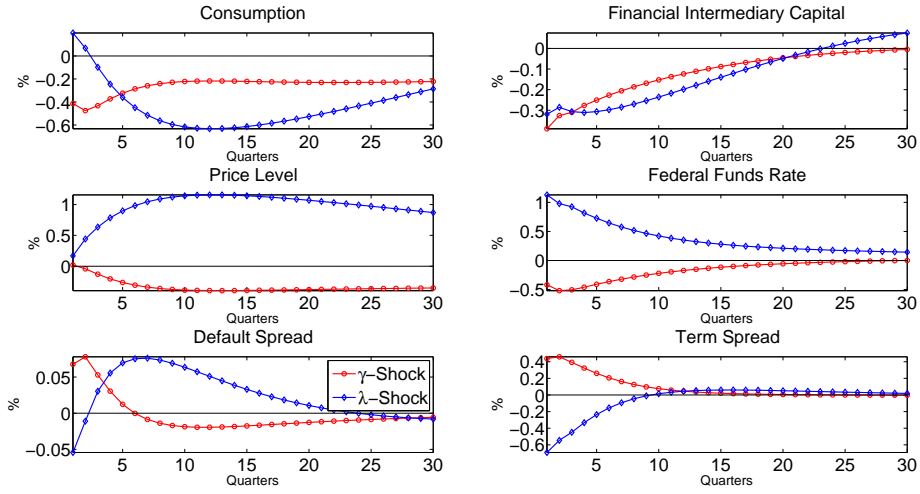
(b) The Contribution of the λ-shock and γ-shock to Capital Ratio Fluctuations

	Intermediary Capital Ratio				Consumption			
	λ-Shock	γ-Shock	λ&γ	Unexpl.	λ-Shock	γ-Shock	λ&γ	Unexpl.
1Q	31.4	34.9	66.3	33.7	4.0	55.5	59.5	40.5
2Q	33.0	34.6	67.6	32.4	2.1	59.7	61.8	38.2
3Q	33.9	33.0	66.9	33.1	5.6	58.3	63.9	36.1
4Q	34.7	31.2	65.9	34.1	12.5	53.7	66.3	33.7
8Q	35.6	26.1	61.7	38.3	41.9	34.6	76.5	23.5
16Q	34.1	21.9	56.0	44.0	61.1	21.3	82.4	17.6
32Q	32.4	20.2	52.6	47.4	60.2	19.5	79.7	20.3

Notes: In panel a, the vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The VAR(2) is estimated on a subsample 1970Q1-2015Q3 given the availability of the intermediary capital ratio series of He, Kelly, and Manela (2017) that starts in 1970. I use the FF55 to construct the λ-shock, and use the CRSP based aggregate SP500 stock market return (over the corresponding T-bill rate) to construct the γ-shock. The table in panel b shows the % fraction of the total forecast error variance that is explained by the λ-shock and the γ-shock over different forecast horizons.

tuations in intermediary balance sheets and their implications for financial markets and the wider economy.

B.9 Pricing the Cross-section of Stock Returns

It is worth noting that the focus of this paper is not the asset pricing performance of the λ-shock. The pricing performance of the given λ-shock can easily be improved by changing the specification of the VAR (e.g. including additional variables such as the excess return on the market[29]). Checking the asset-pricing performance of the λ-shock is therefore only a test as to whether the variables included in the VAR contain information relevant to pricing the given portfolios. Tables 6–9 present the results from the two-pass regression technique of Fama and MacBeth (1973). During this exercise, I treat the uncovered λ-shock as a known factor when estimating the two-pass regression model. To estimate the risk premium associated with the λ-shock, I apply the GMM procedure described in Cochrane (2005) and implemented by Burnside (2011).

Overall, the pricing performance of the VAR (or equivalently, the λ-shock) is comparable with the 3-factor model of Fama and French (1993).[30] Moreover, as explained in the main text (Section 2), finding the λ-shock implies that the other four orthogonalised shocks have zero covariance with the implied SDF, and therefore the associated estimated prices of risk are numerically zero, as shown in panel B of Tables 6–9. Relatedly, the R^2 statistic (computed based on A.15) associated with the one-factor model using the λ-shock is identical to the R^2 for the model using any set of five orthogonalised shocks or in fact the model which uses the five reduced-form VAR residuals.

Moreover, the results are also consistent with Lewellen, Nagel, and Shanken (2010) who pointed out the strong factor structure of the FF25 portfolios which makes it relatively easy to find factors that generate high cross-sectional R^2s. Hence, they prescribed to augment the FF25 with the 30 industry portfolios of Fama-French to relax the tight factor structure of the FF25. Indeed, the cross-sectional R^2 drops drastically from 0.82 to 0.19 for the 1-factor model without a common constant, and it drops from 0.65 to 0.09 for the 3-factor model of Fama-French without a common constant. This can be interpreted as the relevant information content of the VAR being much smaller for pricing the FF55 portfolios than for pricing the FF25 portfolios. Nevertheless, augmenting the VAR to improve pricing performance is unnecessary: the macroeconomic shock that captures all relevant information for pricing the cross section (irrespective of whether the information content is relatively small or large) bears virtually the same economic characteristics as the λ-shock using the FF25 portfolios. The IRFs are similar for the λ-shock using the FF25 and the FF55 (Figures 2 and 17), and the time-series of the shocks implied by the two portfolios have a 0.89 correlation coefficient on the 1964-2015Q3 sample.

---

[29]These results are available upon request.

[30]Applying the 3-factor model to the FF25 portfolios (Table 7) yields similar results to those obtained in the literature (e.g. Petkova (2006)).

Table 6: Results from the Two-pass Regressions, FF55 Portfolios

Factor Prices							R^2
Panel A: 1-factor Model with the λ-Shock							
Constant	λ-Shock						
0.78	0.54						0.28
(0.56) [0.64]	(0.23) [0.26]						
	0.84						0.19
	(0.27) [0.35]						
Panel B: 5-factor Model with the λ- and Other VAR Shocks							
	λ-Shock	Shock2	Shock3	Shock4	Shock5	Shock6	
	0.84	0.00	0.00	0.00	0.00	0.00	0.19
	(0.20) [0.25]						
Panel C: The Fama-French 3-factor Model							
Constant	MKT	HML	SMB				
3.13	-1.44	0.73	0.66				0.40
(0.73) [0.76]	(0.92) [0.95]	(0.42) [0.43]	(0.40) [0.40]				
	1.66	0.75	0.50				0.09
	(0.60) [0.60]	(0.42) [0.43]	(0.40) [0.40]				

Notes: This table reports the cross-sectional regressions using the excess returns on the FF55 portfolios. The coefficients are expressed as percentage per quarter. Panel A presents results for the 1-factor model where the identified λ-shock is used as the sole pricing factor. Panel B presents the results for five-factor model using all orthogonalised shocks from the VAR(2). Panel C presents results for the Fama-French 3-factor model. MKT is the market factor, HML is the book-to-market factor and SMB is the size factor. OLS standard errors are in parentheses, whereas standard errors using the Shanken (1992) procedure are in brackets. The R^2 statistic is computed based on A.15. The sample period is 1964Q1-2015Q3.

Table 7: Results from the Two-pass Regressions, FF25 Portfolios

Factor Prices							R^2
Panel A: 1-factor Model with the λ-Shock							
Constant	λ-Shock						
0.22	1.10						0.83
(0.71) [1.07]	(0.25) [0.38]						
	1.21						0.82
	(0.35) [0.56]						
Panel B: 5-factor Model with the λ- and Other VAR Shocks							
	λ-Shock	Shock2	Shock3	Shock4	Shock5	Shock6	
	1.21	0.00	0.00	0.00	0.00	0.00	0.82
	(0.22) [0.34]						
Panel C: The Fama-French 3-factor Model							
Constant	MKT	HML	SMB				
3.26	-1.62	1.14	0.65				0.73
(0.95) [1.00]	(1.12) [1.16]	(0.41) [0.41]	(0.39) [0.40]				
	1.55	1.21	0.69				0.65
	(0.60) [0.60]	(0.41) [0.42]	(0.39) [0.40]				

Notes: See notes under Table 6.

Table 8: Results from the Two-pass Regressions, 25 Profitability-Size Portfolios

Factor Prices							R^2
Panel A: 1-factor Model with the λ-Shock							
Constant	λ-Shock						
0.03	1.46						0.61
(0.65) [1.17]	(0.49) [0.88]						
	1.48						0.61
	(0.47) [0.86]						
Panel B: 5-factor Model with the λ- and Other VAR Shocks							
	λ-Shock	Shock2	Shock3	Shock4	Shock5	Shock6	
	1.48	0.00	0.00	0.00	0.00	0.00	0.61
	(0.44) [0.79]						
Panel C: The Fama-French 3-factor Model							
	MKT	HML	SMB				
2.78	-1.09	1.03	0.65				0.65
(1.00) [1.04]	(1.17) [1.20]	(0.64) [0.65]	(0.40) [0.40]				
	1.50	1.97	0.63				0.56
	(0.60) [0.60]	(0.73) [0.78]	(0.40) [0.41]				

Notes: See notes under Table 6.

Table 9: Results from the Two-pass Regressions, 25 Investment-Size Portfolios

| Factor Prices							R^2
Panel A: 1-factor Model with the λ-Shock							
Constant	λ-Shock						
0.41	1.01						0.59
(0.61) [0.89]	(0.31) [0.45]						
	1.21						0.56
	(0.37) [0.58]						
Panel B: 5-factor Model with the λ- and Other VAR Shocks							
	λ-Shock	Shock2	Shock3	Shock4	Shock5	Shock6	
	1.21	0.00	0.00	0.00	0.00	0.00	0.56
	(0.20) [0.31]						
Panel C: The Fama-French 3-factor Model							
	MKT	HML	SMB				
1.68	0.06	1.77	0.46				0.76
(1.05) [1.11]	(1.19) [1.24]	(0.52) [0.53]	(0.40) [0.40]				
	1.68	2.10	0.43				0.74
	(0.60) [0.60]	(0.49) [0.51]	(0.40) [0.40]				

Notes: See notes under Table 6.

B.10 The λ-shock, Monetary Policy Shocks and TFP News Shocks

An application of the proposed orthogonalisation strategy to the standard equity portfolios lead to the result that the estimated λ-shock bears a close empirical relationship both with TFP news shocks and with monetary policy shocks. As briefly discussed in the main text, a simple explanation for such an ambiguity is that TFP news shocks and monetary policy shocks are highly correlated in the data.

To provide evidence for this, I use the VAR model of Kurmann and Otrok (2013) to identify a monetary policy shock using Cholesky orthogonalisation as done by Sims (1980), Christiano, Eichenbaum, and Evans (1999) and many others in the monetary policy literature. In this case, I deliberately use exactly the same VAR specification as used by Kurmann and Otrok (2013) when they identified a TFP news shock so that I can learn about differences and similarities across the two identification themes without changing the information set. The upper panel of Figure 19 plots the estimated time-series of the TFP news shocks (black dashed line) against the monetary policy shock series identified with Cholesky orthogonalisation (red solid line). The correlation between the two series is strikingly high (0.96), raising serious questions about the orthogonality of these shocks with respect to one another.

Figure 19: Comparing TFP News Shocks against Monetary Policy Shocks: Results from Kurmann and Otrok (2013)'s VAR and from Smets and Wouters (2007)'s DSGE Model.
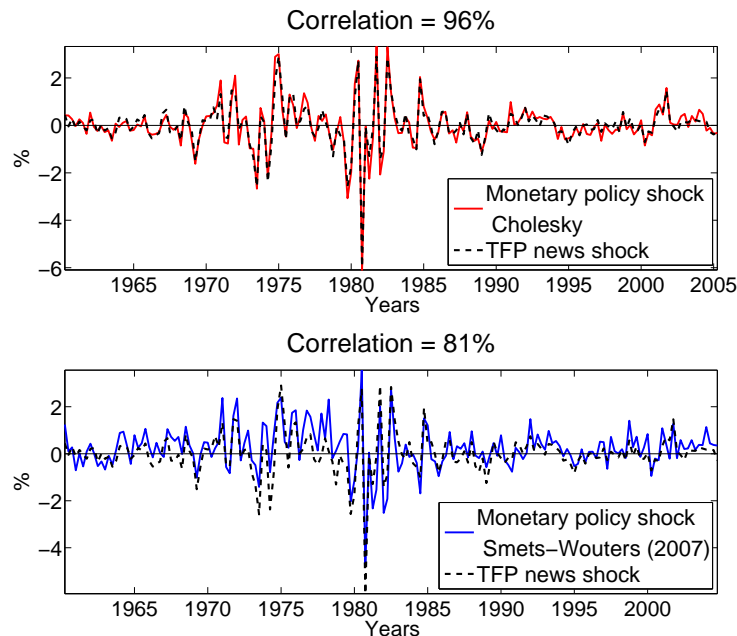
Notes: The TFP news shock series (black dashed line) are the ones plotted in Figure 5 on pp. 2625 of Kurmann and Otrok (2013) who apply the method of Uhlig (2004) to identify a TFP news shock over the period 1959Q2-2005Q2. The monetary policy shock series in the upper panel (red solid line) are identified with Cholesky identification as in Christiano, Eichenbaum, and Evans (1999), using the same variables and lag length as Kurmann and Otrok (2013). The monetary policy shock series in the lower panel (blue solid line) are the estimated time-series of innovations in the Taylor-rule in the DSGE model of Smets and Wouters (2007).

Of course, the identification of monetary policy shocks with Cholesky orthogonalisation is only one of the many possible identification strategies. Therefore, I provide additional evidence from the structural model of Smets and Wouters (2007) which is a dynamic stochastic general equilibrium (DSGE) model estimated with Bayesian methods. Monetary policy shocks in this framework are the estimated innovations in a Taylor-type monetary policy rule. The estimated time-series of these structural innovations from the DSGE model are plotted in the lower panel of Figure 19 (blue solid line) against the TFP news shocks (black dashed line) of Kurmann and Otrok (2013). The correlation between these two series is still remarkably high (0.81).

I interpret these findings that the ambiguous characterisation of the estimated λ-shock does not reflect the weakness of my orthogonalisation theme, but is a result of the high empirical correlation between the two, well-known structural disturbances that the λ-shock resembles. To the best of my knowledge, this empirical regularity has not been documented in the literature yet, and it could be subject to further research. For example, the high empirical correlation may be because of the true correlatedness of these structural disturbances (Curdia and Reis, 2010). An alternative, negative reading of this finding is that it is an identification problem in the literature. To provide some suggestive evidence for this, it is instructive to first review the main assumption of Kurmann and Otrok (2013)'s identification, which builds on the premise that observed technology follows the exogenous process:

$$\log TFP_t = v\left(L\right)\varepsilon_t^{current} + d\left(L\right)\varepsilon_t^{news}, \tag{B.1}$$

which assumes that technology is driven by two uncorrelated innovations: one related to current innovations affecting TFP in t ($\varepsilon_t^{current}$), and the other one (ε_t^{news}) which affects TFP only in $t+1$ onwards. The exogeneity assumption B.1 is used together with zero restrictions on contemporaneous movement on observed TFP. They implement the identification theme following Barsky and Sims (2011) which in turn builds on Uhlig (2004). This entails searching for a structural shock in the VAR which (i) does not move TFP on impact, and (ii) explains the maximal amount of the forecast error variance in TFP over some forecast horizon (40 quarters).

The question then is whether a small amount of violation of assumption B.1 could deliver "TFP news shocks" that can act as monetary policy shocks. I find that assumption B.1 does seem to be violated empirically. For example, the observed utilisation-adjusted TFP measure of Fernald (2012) that Kurmann and Otrok (2013) uses is considerably cyclical in the data. They use vintages of TFP growth that can have about 0.4–0.5 correlations with output growth, compared to most recent vintages that have a lower contemporaneous correlation (Sims, 2016). Of course, correlation coefficients are only crude measures of cyclicality, and it is more instructive to analyse the conditional dynamic relationship in a VAR.

Figure 20: Assessing the Conditional Cyclicality of the Observed TFP measure: The Effect of a Monetary Policy Shock Using the Romer and Romer (2004) Narrative Measure
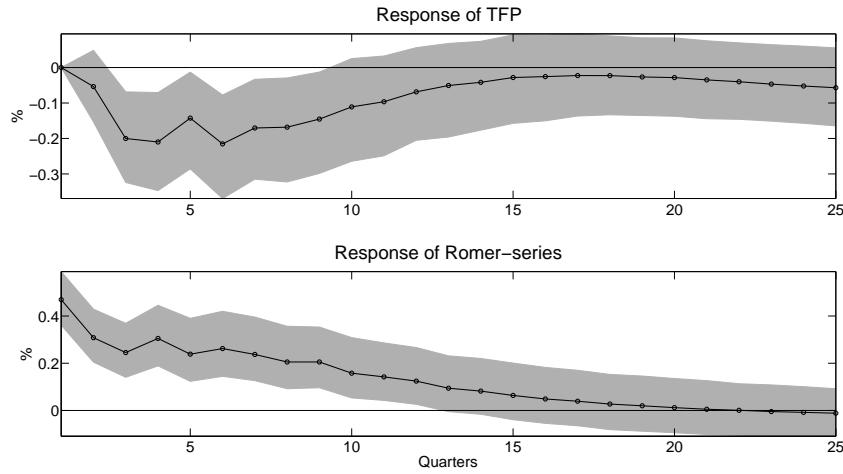
Notes: The Figure shows the IRFs for TFP and the cumulative sum of the Romer and Romer (2004) series (updated by Tenreyro and Thwaites (2016)) from a VAR(4) model. The sample period is 1969Q1-2005Q2. The model includes a constant, but results with a constant and a linear/quadratic trend are very similar. The shaded areas show 95% wild-bootstrapped confidence bands.

Therefore I re-estimate the five-variable VAR model of Kurmann and Otrok (2013) after replacing the short-term interest rate with the cumulative sum of the monetary policy innovations of Romer and Romer (2004) and apply Cholesky orthogonalisation in order to measure the cyclicality of TFP conditional on exogenous monetary policy shocks. Thereby I follow the most recent practice of estimating monetary policy effects in VAR models using narrative measures (Cloyne and Hurtgen, 2016). To focus the attention to the response of TFP, Figure 20 shows only two of the five sets of the IRFs in response to a one standard deviation contractionary shock to monetary policy. Just like in the case of TFP news shock, the monetary policy shock induces a delayed response in TFP. Moreover, the peak response (based on the point estimate) is around 0.2% in absolute value, which is also very similar to the peak effect of a TFP news shock on TFP. The endogenous reaction of TFP to monetary policy shocks displayed by Figure 20 (i) can make it difficult to apply assumption B.1 to identifying a TFP news shock, and as a result (ii) it may be that the 'identified' TFP news shock (ε_t^{news}) is actually picking up some of these monetary policy effects. This could be one of the explanations behind the large empirical correlations displayed by Figure 19.

B.11 Results from the UK

To check whether the results are similar when looking at countries other than the US, I apply the proposed VAR methodology to UK data, covering the period 1970Q1-2012Q4. One advantage of using data for the UK is related to the availability of both comparable monetary policy shock series and comparable test assets across the two countries in question. To estimate the λ-shock, I use the cross-section of 16 equity portfolios (FF16UK), constructed

by Dimson, Nagel, and Quigley (2003). Their portfolio formation closely follows Fama and French (1993), by creating portfolios sorted on size-B/M, whereby breakpoints were applied to the 40th, 60th and 80th percentiles of market capitalisation and to the 25th, 50th and 75th percentiles of book-to-market. To estimate the γ-shock, I use excess returns on the FTSE All-Share index from Chin and Polk (2015), and I also use their series of the Price-Earnings (PE) ratio, as an alternative predictor (given the lack of available CAY measure for the UK).

To keep the empirical model close to the US counterpart presented above, I estimate a VAR(2) model with five macroeconomic variables: log of consumption, log of GDP, log of CPI, the Bank of England policy rate, and the term spread defined as the difference between the ten-year and one-year constant maturity Gilt rates. Given the open-economy nature of the UK and also the lack of available time-series for the default spread, I use the dollar-sterling exchange rate as the sixth variable in the VAR.

Forecasting Excess Returns in the UK As in my baseline model for the US (Table 1), I construct the γ-shock for the UK by maximising the corresponding return forecasting power at 4-quarter horizon, and using the same γ-shock, I compute the results for different horizons ranging from one quarter ahead up to two years ahead. Panel A reports the results using the actual VAR variables as predictors; Panel B shows the results using the Price-Earnings (PE) ratio used by Chin and Polk (2015); Panel C reports the results using the three counterfactual VAR variables induced by the γ-shock; Panel D reports the results using the counterfactual variables induced by all other shocks that are orthogonal to the γ-shock.

Panel A and Panel B of Table 10 are consistent with my baseline results for the US (Table 1) and also corroborate previous evidence for the US on the relevance of valuation ratios to predicting excess results. For example, the last column of the table shows that the PE variable explains around 26% of two-year ahead excess stock market returns; whereas the regression that includes the last three variables of my baseline VAR only explains 16% of excess returns at the same horizon. In contrast, variation in the same VAR variables that is induced by the γ-shock explains about 29% of excess returns at the two-year horizon.

Impulse Response for the UK The upper panel of Figure 21 shows the IRFs for the λ-shock and for Cholesky-orthogonalised interest rate innovations, implied by the UK data. The results are quantitatively very similar to my baseline Figure 21, implied by the US data, with the IRFs of λ-shock being virtually identical to Cholesky interest rate innovations. The lower panel of Figure 21 shows the results for the γ-shock along with the λ-shock. The dynamics are qualitatively very similar to those found the US. An additional finding is that a contractionary λ-shock causes an appreciation of the nominal exchange rate, whereas a negative γ-shock causes a depreciation.

Table 10: Forecasting Excess Returns in the UK

| | \multicolumn{8}{c}{Forecast Horizon H} |
	1Q	2Q	3Q	4Q	5Q	6Q	7Q	8Q
\multicolumn{9}{c}{Model A: Actual VAR Variables}
BoE	0.20	0.44	0.70	0.93	1.19	1.50	1.78	2.11
	(0.84)	(1.03)	(1.26)	(1.51)	(1.73)	(1.84)	(2.03)	(2.30)
TERM	1.07	2.29	3.39	4.23	4.73	5.26	6.03	6.75
	(1.32)	(1.66)	(1.94)	(2.16)	(2.21)	(2.27)	(2.46)	(2.59)
EXCH	-0.06	-0.14	-0.21	-0.27	-0.32	-0.38	-0.45	-0.51
	(-1.41)	(-1.32)	(-1.25)	(-1.15)	(-1.10)	(-1.22)	(-1.43)	(-1.70)
	[0.01]	**[0.04]**	**[0.07]**	**[0.08]**	**[0.09]**	**[0.11]**	**[0.14]**	**[0.16]**
\multicolumn{9}{c}{Model B: PE}
PE	-0.37	-0.79	-1.23	-1.68	-2.10	-2.50	-2.95	-3.30
	(-1.33)	(-1.58)	(-2.13)	(-2.99)	(-3.80)	(-3.80)	(-3.69)	(-3.77)
	[0.02]	**[0.04]**	**[0.07]**	**[0.11]**	**[0.14]**	**[0.18]**	**[0.22]**	**[0.26]**
\multicolumn{9}{c}{Model C: Counterfactual VAR Variables Induced by the γ-Shock}
BoE	0.18	0.31	0.47	0.61	0.75	0.94	1.14	1.35
	(0.76)	(0.69)	(0.75)	(0.79)	(0.84)	(0.95)	(1.10)	(1.30)
TERM	1.28	2.75	4.20	5.59	6.70	7.83	8.97	10.02
	(1.13)	(1.43)	(1.83)	(2.27)	(2.52)	(2.70)	(2.98)	(3.22)
EXCH	-0.12	-0.24	-0.35	-0.46	-0.56	-0.67	-0.78	-0.88
	(-2.00)	(-1.53)	(-1.32)	(-1.23)	(-1.23)	(-1.36)	(-1.58)	(-2.03)
	[0.02]	**[0.06]**	**[0.09]**	**[0.13]**	**[0.16]**	**[0.20]**	**[0.24]**	**[0.29]**
\multicolumn{9}{c}{Model D: Counterfactual VAR Variables Induced by All Other Shocks}
BoE	0.11	0.24	0.39	0.49	0.63	0.86	1.08	1.39
	(0.50)	(0.58)	(0.69)	(0.72)	(0.81)	(0.93)	(1.05)	(1.23)
TERM	0.23	0.39	0.47	0.00	-0.74	-1.14	-0.93	-0.64
	(0.31)	(0.29)	(0.27)	(0.00)	(-0.26)	(-0.31)	(-0.24)	(-0.16)
EXCH	-0.02	-0.04	-0.07	-0.07	-0.06	-0.09	-0.13	-0.17
	(-0.29)	(-0.34)	(-0.38)	(-0.29)	(-0.21)	(-0.29)	(-0.41)	(-0.52)
	[-0.02]	**[-0.01]**	**[-0.01]**	**[-0.01]**	**[-0.00]**	**[0.01]**	**[0.02]**	**[0.03]**

Notes: The table reports results from regressions of excess returns on lagged variables. H denotes the return horizon in quarters. The dependent variable is the sum of H log excess returns on the CRSP based S&P Composite Index. The regressors are one-period lagged values of actual time-series of the Bank of England base rate (BoE), the term-spread (TERM) and the dollar-sterling nominal exchange rate (EXCH) in Model A, the Price-Earnings (PE) ratio of Chin and Polk (2015) in Model B, and the counterfactual time-series of BoE, TERM and EXCH (induced by the γ-shock) from a six-variable VAR(2) estimated over 1970Q1-2012Q4. The γ-shock is constructed so that the corresponding forecast power at the four-quarter horizon is maximised. For each of the three regressions, the table reports the OLS estimates of the regressors, the t-statistics using the Hansen and Hodrick (1980) correction (as implemented in Cochrane (2011)) are in parentheses, and adjusted R^2 statistics are in the bolded square brackets. Both the PE measure and the counterfactual predictors are treated as known variables.

Moreover, similar to the US case, the estimated time-series of the λ-shock is empirically related to monetary policy shocks. The monetary policy shock series of Cloyne and Hurtgen (2016) and the estimated λ-shock series have around 60% correlation on the overlapping sample (1975Q1-2007Q4).[31] Overall, the results obtained for the UK are similar to those obtained for the US.

---

[31] The methodology of Cloyne and Hurtgen (2016) follows that of Romer and Romer (2004) by trying to eliminate much of the endogenous movement between the interest rate and other macroeconomic variables as well as to control for the effects related to current expectations of future economic conditions.

Figure 21: Results from the UK

(a) Impulse Responses to a λ-shock and to an Interest Rate Shock
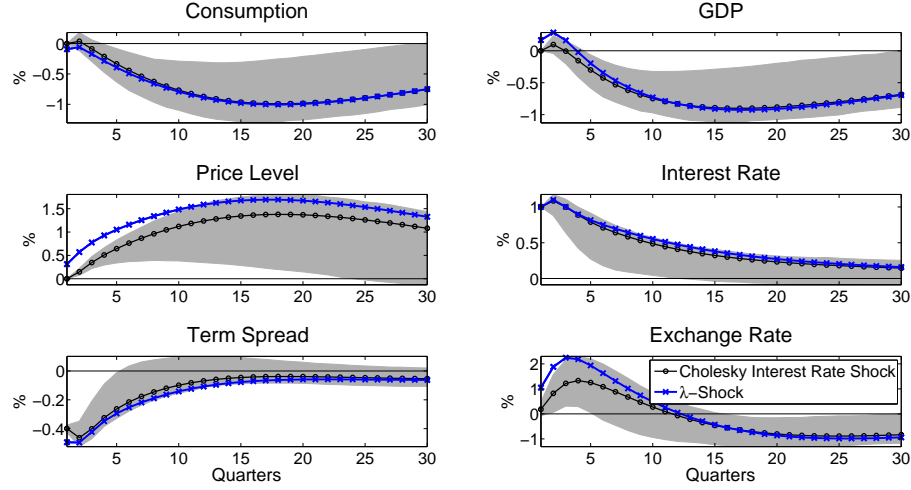
(b) Impulse Responses to a γ-Shock and to a λ-Shock
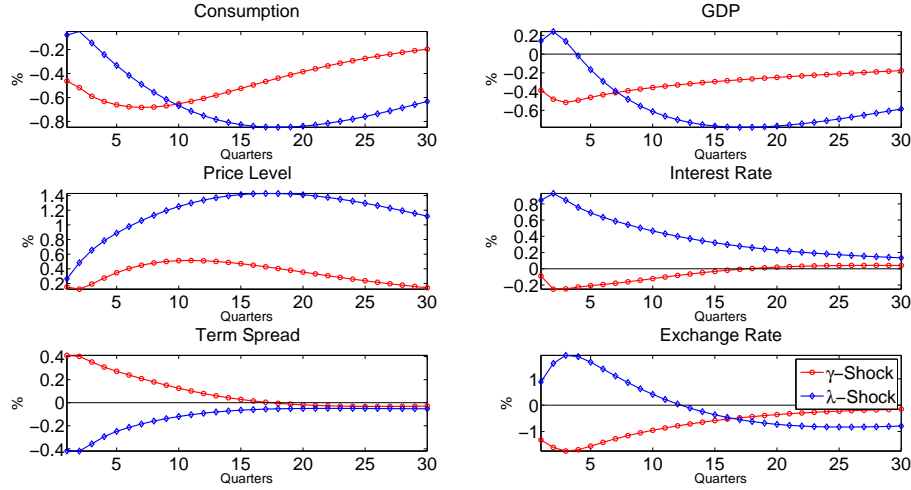
Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. The VAR(2) is estimated on the sample 1970Q1-2012Q4. The FF16UK is (Dimson, Nagel, and Quigley, 2003) to construct the λ-shock. While the VAR is estimated on full sample, the rotation of the variance-covariance matrix is based on the 1970Q1-2001Q4, because the FF16UK series end in 2001Q4. The excess returns on the FTSE All-Share index (Chin and Polk, 2015) are used to constructed the γ-shock. In the upper panel, the blue crossed lines are λ-shock, and the blacked circles lines are Cholesky-orthogonalised interest rate shock with the associated 95% confidence band (using wild-bootstrap). In the upper panel, the IRFs are normalised to increase the interest rate by 100bp. In the lower panel, the magnitude of both shocks is one standard deviation.